IBM DS8000 Series
Version 9 Release 0

Command-Line Interface User's Guide



SC27-9562-00

Note
Before using this information and the product it supports, read the information in the "Notices" on page 553 section.

This edition applies to version 8, release 5, Modification 4 of the *IBM DS8000 Series Command-Line Interface User's Guide* and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 2004, 2019.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this book v
Who should use this book v
Syntax diagram conventions v
Publications and related information v
IBM Publications Center viii
Sending comments ix

Summary of Changes xi

Chapter 1. Installing, upgrading, and uninstalling the DS CLI 1
Operating systems that support the DS CLI 1
Installing the DS CLI 1
Requirements 1
Overview of installation methods 3
Correcting the Java virtual machine not found error 6
Upgrading the DS CLI on your system 7
Uninstalling the DS CLI 8
Overview of uninstallation methods 8
Running and configuring the DS CLI 9
Creating a default CLI profile 10
Setting up user accounts by using the DS CLI . . 17
Activating your machine and feature licenses by using the DS CLI 22

Chapter 2. Running the DS CLI 25
Logging in to the DS CLI 25
Using the DS CLI single-shot command mode . . . 26
Using the DS CLI script command mode 27
Using the DS CLI interactive command mode (history and reports) 28
Obtaining the serial (storage image ID) number by using the DS CLI 29
DS CLI command help 29
Obtaining and interpreting DS CLI exit codes . . . 30
DS CLI operational limitations 32
Messages in the CLI and management console server 33

Chapter 3. CLI commands 35
About CLI commands 35
Understanding the syntax diagrams 35
Common command flags 37
Application key commands 38
applykey 38
lskey 39
Cloud commands 40
lscloudserver 41
managecloudserver 42
mkcloudserver 43
rmcloudserver 46
Cluster commands 46
chcluster 46
lscluster 48
mkcluster 50
rmcluster 50
showcluster 51
Copy Services commands 53
FlashCopy commands 53
Remote FlashCopy commands 76
Remote Mirror and Copy path commands . . . 95
Remote Mirror and Copy commands 108
Global Mirror commands 137
Global Mirror session commands 153
Framework commands 160
dscli 161
echo 162
exit 163
help 163
helpmsg 164
quit 165
setenv 165
setoutput 167
showenv 170
ver 171
Install software commands 173
installsoftware 173
lssoftware 174
I/O port and host connection configuration commands 175
Storage image I/O port commands 175
Host commands 199
Host port commands 226
chhostport 226
lshostport 227
mkhostport 231
rmhostport 231
showhostport 232
chzhyperlink 234
lszhyperlink 235
lszhyperlinkport 236
showzhyperlinkport 237
Network port commands 238
setnetworkport 239
lsnetworkport 240
shownetworkport 243
Offload file commands 245
offloadauditlog 245
offloadfile 248
Physical resource information commands 252
lsda 252
lsddm 255
lsframe 260
lshba 261
lsstgencl 262
Remote event notification commands 265
chsyslogserver 265
lssyslogserver 267
mksyslogserver 268

© Copyright IBM Corp. 2004, 2019

rmsyslogserver 270
Resource Group commands 271
chresgrp 271
lsresgrp 272
manageresgrp 274
mkresgrp 277
rmresgrp 278
showresgrp 278
Security commands 280
chauthpol 280
chpass 281
chuser 282
cpauthpol 285
lsauthpol 285
lsuser 286
lsusergroup 288
lswuiuser 289
managepwfile 290
mkauthpol 292
mkuser 293
rmauthpol 295
rmuser 295
setauthpol 296
showauthpol 305
showpass 310
showuser 310
testauthpol 312
who 312
whoami 315
Data encryption and security commands . . . 316
User access security commands 351
Storage complex configuration commands 358
chsp 359
setvpn 361
lsvpn 362
showsp 363
Storage configuration commands 365
Array site specific commands 365
Array specific commands 371
Rank specific commands 379
Extent pool specific commands 392
Address group specific commands 409
Logical control unit specific commands 410
CKD logical volume specific commands . . . 420
Logical subsystem specific commands 455
Fixed block logical volume specific commands 461
Volume group specific commands 499
Advanced operation commands 509
Space-efficient storage commands 512
Storage image configuration commands 533
chsi 533
diagsi 536
lsserver 537
lssi 539
showsi 541
Storage system configuration commands 545
chsu 545
lssu 547
showsu 549

Notices 553
Trademarks 554

Index 555

About this book

This book describes information about the DS8000® series command-line interface. The first chapter provides an overview of the DS8000 series. Subsequent chapters describe installing, upgrading, removing, and running the DS CLI to configure and run your storage systems.

Who should use this book

This book is intended for system administrators or others who use the DS CLI to install and manage the IBM® DS8000 systems.

Syntax diagram conventions

A syntax diagram uses symbols to represent the elements of a command and to specify the rules for using these elements.

The following table displays the conventions that are used in the command syntax for the product.

Table 1. Command syntax conventions

Syntax convention	Description	Example
Command	A command is the first word or set of consecutive characters.	**help**
Option	An option is a character, set of consecutive characters, or a word that follows the command and any arguments.	[on]
Variable	A variable is any set of consecutive characters or word that follows an option. Variables are in *italic* typeface and can use capitalized letters in the character string to aid in reading comprehension.	*timeout_in_sec*
Vertical bar (|)	Mutually exclusive options are separated by a vertical bar (|).	[on | off]
Ellipsis (...)	An ellipsis (...) indicates that the previous option can be repeated multiple times with different values. It can be used inside or outside of brackets.	source1:target1 [,source2:target2][,...]

Publications and related information

Product guides, other IBM publications, and websites contain information that relates to the IBM DS8000 series.

To view a PDF file, you need Adobe Reader. You can download it at no charge from the Adobe website(get.adobe.com/reader/).

Online documentation

The IBM DS8000 series online product documentation (http://www.ibm.com/support/knowledgecenter/ST5GLJ_8.1.0/com.ibm.storage.ssic.help.doc/f2c_securitybp.html) contains all of the

© Copyright IBM Corp. 2004, 2019

information that is required to install, configure, and manage DS8000 storage systems. The online documentation is updated between product releases to provide the most current documentation.

Publications

You can order or download individual publications (including previous versions) that have an order number from the IBM Publications Center website(https://www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss). Publications without an order number are available on the documentation CD or can be downloaded here.

Table 2. DS8000 series product publications

Title	Description	Order number
IBM DS8900F Introduction and Planning Guide	This publication provides an overview of the new DS8900F, the latest storage system in the DS8000 series. The DS8900F provides two system types: DS8910F Flexibility Class models 993 and 994 and DS8950F Agility Class models 996 and E96.	V9.0.0 SC27-9560-00
IBM DS8882F Introduction and Planning Guide	This publication provides an overview of the DS8882F, the latest storage system in the DS8000 series. The DS8882F provides the new model 983. This publication provides an overview of the product and technical concepts for DS8882F.	V8.5.3 GC27-9259-02 V8.5.1 GC27-9259-01 V8.5.0 GC27-9259-00
IBM DS8880 Introduction and Planning Guide	This publication provides an overview of the product and technical concepts for DS8880. It also describes the ordering features and how to plan for an installation and initial configuration of the storage system.	V8.5.3 GC27-8525-18 V8.5.1 GC27-8525-17 V8.5.0 GC27-8525-16 V8.4.0 GC27-8525-15 V8.3.3 GC27-8525-14 V8.3.1 GC27-8525-13 V8.3.0 GC27-8525-12 V8.2.3 GC27-8525-11 V8.2.1 GC27-8525-09 V8.2.0 GC27-8525-07 V8.1.1 GC27-8525-06 V8.1.0 GC27-8525-05 V8.0.1 GC27-8525-04 GC27-8525-03 V8.0.0 GC27-8525-02
IBM DS8870 Introduction and Planning Guide	This publication provides an overview of the product and technical concepts for DS8870. It also describes the ordering features and how to plan for an installation and initial configuration of the storage system.	V7.5.0 GC27-4209-11 V7.4.0 GC27-4209-10 V7.3.0 GC27-4209-09 V7.2.0 GC27-4209-08 V7.1.0 GC27-4209-05 V7.0.0 GC27-4209-02

Table 2. DS8000 series product publications (continued)

Title	Description	Order number
IBM DS8000 Command-Line Interface User's Guide	This publication describes how to use the DS8000 command-line interface (DS CLI) to manage DS8000 configuration and Copy Services relationships, and write customized scripts for a host system. It also includes a complete list of CLI commands with descriptions and example usage.	V9.0.0 SC27-9562-00 V8.5.4 SC27-8526-11 V8.5.1 SC27-8526-10 V8.5.0 SC27-8526-09 V8.3.3 SC27-8526-08 V8.3.1 SC27-8526-07 V8.3.0 SC27-8526-06 V8.2.3 SC27-8526-05 V8.2.2 SC27-8526-04 V8.2.0 SC27-8526-03 V8.1.1 SC27-8526-02 V8.1.0 SC27-8526-01 V8.0.0 SC27-8526-00 V7.5.0 GC27-4212-06 V7.4.0 GC27-4212-04 V7.3.0 GC27-4212-03 V7.2.0 GC27-4212-02 V7.1.0 GC27-4212-01 V7.0.0 GC27-4212-00 V6.3.0 GC53-1127-07
IBM DS8000 Host Systems Attachment Guide	This publication provides information about attaching hosts to the storage system. You can use various host attachments to consolidate storage capacity and workloads for open systems and IBM Z hosts.	V9.0.0 SC27-9563-00 V8.0.0 SC27-8527-00 V7.5.0 GC27-4210-04 V7.4.0 GC27-4210-03 V7.2.0 GC27-4210-02 V7.1.0 GC27-4210-01 V7.0.0 GC27-4210-00 V6.3.0 GC27-2298-02
IBM DS8800/DS8700 RESTful API Guide	This publication provides an overview of the Representational State Transfer (RESTful) API, which provides a platform independent means by which to initiate create, read, update, and delete operations in the DS8000 and supporting storage devices.	V1.3 SC27-9235-00 V1.2 SC27-8502-02 V1.1 SC27-8502-01 V1.0 SC27-8502-00

Table 3. DS8000 series warranty, notices, and licensing publications

Title	Location
IBM Warranty Information for DS8000 series	IBM Support Portal website
IBM Safety Notices	IBM Systems Safety Notices
IBM Systems Environmental Notices	IBM Support Portal website
International Agreement for Acquisition of Software Maintenance (Not all software will offer Software Maintenance under this agreement.)	IBM Support Portal website
IBM License Agreement for Machine Code	IBM Support Portal website

See the Agreements and License Information CD that was included with the DS8000 series for the following documents:
- License Information
- Notices and Information
- Supplemental Notices and Information

Related websites

View the websites in the following table to get more information about DS8000 series.

Table 4. DS8000 series related websites

Title	Description
IBM website (ibm.com®)	Find more information about IBM products and services.
IBM Support Portal website	Find support-related information such as downloads, documentation, troubleshooting, and service requests and PMRs.
IBM Directory of Worldwide Contacts website(www.ibm.com/planetwide)	Find contact information for general inquiries, technical support, and hardware and software support by country.
IBM DS8000 series website (www.ibm.com/servers/storage/disk/ds8000)	Find product overviews, details, resources, and reviews for the DS8000 series.
IBM Redbooks® (www.redbooks.ibm.com/)	Find technical information developed and published by IBM International Technical Support Organization (ITSO).
IBM System Storage® Interoperation Center (SSIC) website (www.ibm.com/systems/support/storage/config/ssic)	Find information about host system models, operating systems, adapters, and switches that are supported by the DS8000 series.
IBM Data storage feature activation (DSFA) website (www.ibm.com/storage/dsfa)	Download licensed machine code (LMC) feature keys that you ordered for your DS8000 storage systems.
IBM Fix Central (www-933.ibm.com/support/fixcentral)	Download utilities such as the IBM Easy Tier® Heat Map Transfer utility and Storage Tier Advisor tool.
IBM Java™ SE (JRE)(www.ibm.com/developerworks/java/jdk)	Download IBM versions of the Java SE Runtime Environment (JRE), which is often required for IBM products.
IBM Security Key Lifecycle Manager online product documentation(www.ibm.com/support/knowledgecenter/SSWPVP/)	This online documentation provides information about IBM Security Key Lifecycle Manager, which you can use to manage encryption keys and certificates.
IBM Spectrum Control online product documentation in IBM Knowledge Center (www.ibm.com/support/knowledgecenter)	This online documentation provides information about IBM Spectrum Control, which you can use to centralize, automate, and simplify the management of complex and heterogeneous storage environments including DS8000 storage systems and other components of your data storage infrastructure.
DS8880 Code Bundle Information website(www.ibm.com/support/docview.wss?uid=ssg1S1005392)	Find information about code bundles for DS8880. The version of the currently active installed code bundle displays with the DS CLI `ver` command when you specify the -l parameter.

IBM Publications Center

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

Procedure

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

Sending 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 DS8000 series documentation:

Send your comments by email to ibmkc@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

Summary of Changes

DS8000 series Version 9, Release 0 introduces the following enhancements. For DS8000 series information, see the *IBM DS8000 Introduction and Planning Guide* for your specific model.

Version 9, Release 0

This table provides the current technical changes to the DS8000 series.

Function	Description
Modified commands	**chfbvol** command, volume expansion for IBM i volumes is supported. A message displays before proceeding with the volume expansion. For information, see the "chfbvol" on page 461 command. The following commands include support for certificate and security enhancements, including support for the IBM Fibre Channel Endpoint Security feature. • **lsioport** and **setioport** commands, the **-security** parameter displays the current security setting of an I/O port. The values include disabled (the default), enabled, and enforced. See the "lsioport" on page 175 and "setioport" on page 184 commands for more information. • **showkeygrp** command, example output has been enhanced to include IBM Fibre Channel Endpoint Security information as well as certificate information. For information, see the "showkeygrp" on page 340 command. • **lskeygrp** and **mkkeygrp** commands display the type of encryption that is used by the key group to support IBM Fibre Channel Endpoint Security. For information, see the "lskeygrp" on page 317 and "mkkeygrp" on page 334 commands. • **lskeymgr**, **mkkeymgr**, and **managekeymgr** commands display the type of encryption that is used by the key manager to support IBM Fibre Channel Endpoint Security. See the "lskeymgr" on page 322, "lskeymgr" on page 322, and "managekeymgr" on page 330 commands for more information.

© Copyright IBM Corp. 2004, 2019

Chapter 1. Installing, upgrading, and uninstalling the DS CLI

Before you decide to install the DS CLI on your system, familiarize yourself with the operating systems that support this application. Consider the tasks that are involved in upgrading your system. Keep in mind the operational limitations that are associated with the DS CLI.

Note: You can install the DS CLI using the DS CLI installation CD by using the procedures described in the following sections. The DS8000 Storage Management GUI enhancement features an Embedded DS CLI window that you can use to conveniently enter DS CLI commands. You can use the window to enter commands or scripts for the DS CLI. You do not have to install the DS CLI to access and process DS CLI commands.

The DS CLI installation CD comes with the release bundle documentation for IBM. The installation CD ISO image files are also available online. See "Release notes" documented in the DS8000 IBM Knowledge Center website.

Operating systems that support the DS CLI

The DS command-line interface (CLI) can be installed on a variety of operating systems. Refer to the list of operating systems to ensure that your operating system software and its version can support the installation of the DS CLI.

You can install the DS CLI on machines that use one of the following operating systems:

- AIX® 7.1, 7.2
- Linux, Red Hat Enterprise Linux [RHEL] 7
- Linux, SUSE Linux, Enterprise Server (SLES) 12
- VMware ESX 6 Console
- IBM i 7.2
- Oracle Solaris 11
- Microsoft Windows Server 2012 and Windows 8, 8.1, 10

Installing the DS CLI

On most systems you can install the DS CLI using a silent mode, console mode, or by using a GUI application mode.

Requirements

The DS CLI can be used by open systems hosts to start and manage FlashCopy® and Metro and Global Mirror functions through batch processes and scripts. This information provides key considerations for a DS CLI installation on various supported operating systems.

General preinstallation specifics for supported operating systems

The following list provides information for installing the DS CLI on one of the supported operating systems. This information includes the location of the installers for each supported operating system. The installers are installed in the images\HMC\Disk1\InstData directory, and sorted into folders by operating system.

- The following table provides the installation file location, by operating system. See the additional instructions that follow this table.

© Copyright IBM Corp. 2004, 2019

Table 5. Installation file location by operating system

Supported host systems	Installation file location
IBM AIX 7.1, 7.2	images/HMC\Disk1\InstData\AIX\NoVM\ dsclisetup.bin
Linux Red Hat Enterprise Server, RHEL 7	images/HMC\Disk1\InstData\Linux\NoVM\ dsclisetup.bin
Linux, SUSE Linux Enterprise Server (SLES) 12	images/HMC\Disk1\InstData\Linux\NoVM\ dsclisetup.bin
Oracle Solaris 11	images/HMC\Disk1\InstData\Solaris\NoVM\ dsclisetup.bin
IBM i 7.2	images/HMC\Disk1\InstData\Windows\NoVM\ dsclisetup.exe
VMware ESX 6 Console	images/HMC\Disk1\InstData\Linux\NoVM\ dsclisetup.bin
Microsoft Windows Server 2012, Windows, and Windows 8, 8.1, 10	images/HMC\Disk1\InstData\Windows\NoVM\ dsclisetup.exe

- Java Platform (Java 8 or later) must be installed on your system. The installation program checks for this requirement during installation and does not install the DS CLI if you do not have Java Platform, Enterprise Edition (Java 8 or later).
- AIX installation considerations.
  - The LIBPATH environment variable can interfere with the installation of the DS CLI and can result in the display of the Java Virtual Machine Not Found Error. To avoid this interference, you must disable the LIBPATH environment variable before you install the DS CLI. After the installation of the DS CLI, you must enable the LIBPATH environment variable so that it can be used with other applications.
  - Run the following commands to sequentially disable the LIBPATH environment variable, install the DS CLI, and restore the LIBPATH environment variable:

```
export LIBSAVE=$LIBPATH
  unset LIBPATH
  AIX/NoVM/dsclisetup.bin LAX_VM /opt/ibm-Java-whatever/java/bin/java
  export LIBPATH=$LIBSAVE
  unset LIBSAVE
```

- Windows installation considerations.
  - The User Access Control (UAC) settings for Windows 8 and later, or Windows Server 2012 and later, might not allow for exporting files (by using the **offloadfile** command) to a directory that requires elevated privileges. Unfortunately, the Windows operating system returns success in exporting the files and the **offloadfile** command, but the files do not exist in the specified directory. To work around this problem, complete the following steps:
    - Select a different directory that does not require elevated privileges to create a file.
    - Right-click the DSCLI desktop shortcut and select **Run as Administrator**.
- IBM i model installation considerations.

  Note: The installation of DS CLI on an IBM i model is done remotely from a Windows platform. You cannot run the DS CLI installer directly on an IBM i model.

  The IBM i model and i5/OS must meet the following requirements before the DS CLI can be installed:
  - Prerequisites include
    - The latest Java group program temporary fixes (PTF)
    - i5/OS 5722-SS1 option 34 - Digital certificate manager
  - If you are installing onto an IBM i model, ensure that the workstation from which you are installing is network-attached to the iSeries server.

  - During the installation of the DS CLI onto an IBM i model, provide the following information:
    - The name of the IBM i server to which you are installing the DS CLI.
    - The user name and password that are used to access the designated IBM i server.
  - The IBM i TCP attributes for FTP must be set to the library format (the default) when you install the DS CLI. Use the following steps to ensure that the files can be moved to the IBM i partition during the installation:
    - Collect the current attributes for NAMEFMT, CURDIR (listed on the CHGFTPA command).
    - Enter the following command: CHGFTPA NAMEFMT(*LIB) CURDIR(*CURLIB)
    - Restart the FTP server on the IBM i partition.
    - Install the DS CLI.
    - Restore the attributes for NAMEFMT, CURDIR collected in step 1 with the CHGFTPA command.
    - Restart the FTP server again.
  - When you install the DS CLI onto an IBM i model, a `_uninst` folder is created on the Windows desktop. Save this folder for uninstallation in the future.
- The installation process installs the DS CLI in the following default directories:

  AIX /opt/ibm/dscli

  Linux /opt/ibm/dscli

  Oracle Solaris
  /opt/ibm/dscli

  Windows, 32-bit system
  C:\Program Files\IBM\dscli

  Windows, 64-bit system
  C:\Program Files (x86)\IBM\dscli

  IBM i /ibm/dscli

  VMware
  /opt/ibm/dscli

  z/OS® /opt/IBM/CSMDSCLI/dscli
- z/OS is supported on CSM Version 6.1.4 and later. You can install the DS CLI along with Copy Services Manager on z/OS installations. The DS CLI runs under UNIX Services for z/OS and has a separate function modification identifier (FMID) HIWN61K.

  You can also install the DS CLI separately from Copy Services Manager. For more information, see the *IBM DSCLI on z/OS Program Directory* for more information. Use the order number (GI13-3563) to search for the publication from the IBM Publications Center website (www.ibm.com/shop/publications/order/).
- Regardless of the operating system and DS series that you use, activate your license activation codes (part of the DS Storage Management GUI postinstallation instructions). Then, you can use the CLI commands that are associated with Copy Services functions.

Overview of installation methods

The following information provides an overview of the installation methods that are supported and how to access them.

You can install the DS CLI by using one of several methods.

- Graphical mode
- Console mode
- Unattended (silent mode)

Requirements and restrictions

Note the following requirements before you install the DS CLI:

- You must have a version of Java 1.8 or later that is installed on your system in a standard directory. The DS CLI installer checks the standard directories to determine whether a version of Java 1.8 or later exists on your system. If this version is not found in the standard directories, the installation fails.

  Note: Ensure that you use Java 1.8 or later.
- If the DS CLI is installed on your client or host system, you must end any active sessions of the DS CLI before you run the DS CLI installation CD.
- IBM and i5/OS installations have the following requirements:
  - The latest Java group program temporary fixes (PTF)
  - The i5/OS 5722-SS1 option 34 - Digital certificate manager
  - The workstation that you are installing from must be connected to the i5/OS through an IP connection.
- You can use the *-i* parameter to specify any user interface mode when you install the DS CLI: **-i [gui | console | silent]**. You do not have to specify the mode in the command unless you want to use something other than the default mode.
  - The default mode for installing Windows is **gui**.
  - The default for UNIX and Linux is **console** mode.
- While in console mode, you can type `back` to return to the previous screen, or `quit` to exit the installation.

Graphical mode

To install the DS CLI by using the graphical mode, start the setup file that is appropriate for your operating system and proceed through the installation panels.

You can find the setup file on the installation CD by navigating to images/HMC\Disk1\InstData, and then selecting your platform to find the appropriate setup file. For example, in Windows, the path would be images/HMC\Disk1\InstData\Windows\NoVM\dsclisetup.exe.

Console mode

The console mode is primarily used for installations on a Linux operating system or on a UNIX operating system without an X display. You can run the installer from a command prompt on a Windows operating system.

Complete the following steps to install the DS CLI by using the console mode:

1. Insert the DS CLI Installation CD into the CD drive.
2. Open a command prompt and navigate to the location of the `dsclisetup` file on the DS CLI CD. You can find the setup file by navigating to `images/HMC\Disk1\InstData`, and then selecting your platform to find the appropriate setup file. For example, in Windows the path would be `images/HMC\Disk1\InstData\Windows\NoVM\dsclisetup.exe`.
3. Type the following on the command line: **dsclisetup<.exe | .bin> -i console**
   - For Windows, type: `dsclisetup.exe -i console`
   - For Linux, type: `dsclisetup.bin -i console`
   - For an installation onto an IBM i system from a Windows operating system, type: `setupwin32console.exe -os400`
4. Proceed through the rest of the installation windows.

Unattended (silent mode)

Using the unattended (silent) mode of installation, you can install the DS CLI from the command line using default selections without prompts or feedback. You can create a configuration file and use a text editor to change the default installation selections.

To install the DS CLI in unattended (silent mode), complete the following steps:

1. Log on to your system as an administrator.
2. Insert the DS CLI installation CD into the CD drive.
3. Open a command prompt and navigate to the location of the dsclisetup file on the DS CLI CD. You can find the setup file by navigating to *images/HMC\Disk1\InstData*, and then selecting your platform to find the appropriate setup file. For example, in Windows the path would be *images/HMC\Disk1\InstData\Windows\NoVM\dsclisetup.exe*.
4. Issue the following command at the command prompt: **dsclisetup.exe -i silent**. Press the **Enter** key on your keyboard to start the installation process in unattended (silent) mode.

   The silent installation process applies all the default options to your installation.If you want to modify the default options, go to the next step.

   Note: Initially the DS CLI installer checks your standard directories for the correct version of Java. If the correct version of Java is not found, you receive the following message:

   - If you are using Windows, the following message is displayed:

     ```
     LaunchAnywhere Error: Could not find a valid Java virtual
     machine to load.

     You may need to reinstall a supported Java virtual machine.
     ```

   - If you are using UNIX or Linux, the following message is displayed:

     ```
     No Java virtual machine could be found from your PATH
     environment variable.  You must install a VM prior to
     running this program.
     ```

   If you receive this message, see "Correcting the Java virtual machine not found error" on page 6.
5. Optionally, you can generate a configuration file in a non-silent mode, then use this file in subsequent silent installs. For example:
   a. Execute the following command to install a sample instance of dscli. **dsclisetup.exe -r "c:\install.properties"** The resulting file, *install.properties*, will contain all of the install settings.
   b. Modify the settings in *install.properties* if needed.
   c. You can use this generated configuration file in silent installs with the following command: **dsclisetup.exe -i silent -f "c:\install.properties"**

Post-installation notes

Notes:

1. The DS CLI is installed in the following two places in i5/OS:
   - IFS directory *IBM/DS_CLI*. This directory contains the profiles, .EXE files, Java .JAR files, readme files, and so on.
   - The QDSCLI library. This library contains executable code.

     Note: Beginning in Version 6 Release 1, the QDSCLI library output is stored in the output file that you specified in the DS CLI. Errors will be stored in a file of the same name with `.error` appended to it. If you programmed anything using the output, update your programs to look in both the specified file and the `.error` file to get a complete view of the results.
2. Before you can start the DS CLI from the i5/OS, you must add the QDSCLI library to the i5/OS library list.

3. You can check the following directories to verify that the DS CLI has been installed for your operating system:

   AIX /optibm/dscli

   Linux /optibm/dscli

   Oracle Solaris
   /optibm/dscli

   Windows
   (32-bit) C:\Program FilesIBM\dscli

   (64-bit) C:\Program Files (x86)IBM\dscli

   IBM i /ibm/dscli

   z/OS /opt/IBM/CSMDSCLI/dscli

Correcting the Java virtual machine not found error

Complete this task to correct the `Java Virtual Machine Not Found` error.

About this task

The `Java Virtual Machine Not Found` error occurs when the DS CLI installer cannot find the correct version of Java in the standard directories of your system. You must have Java 8 or later on your system for the DS CLI to work.

Notes:

1. This error might also occur if you are installing the DS CLI on an AIX system. The LIBPATH environment variable can interfere with the installation of the DS CLI and can result in the display of the `Java Virtual Machine Not Found` error. To prevent this error, disable the LIBPATH environment variable before you install the DS CLI. After the installation of the DS CLI, enable the LIBPATH environment variable so that it can be used with other applications.

If Java 8 or later is not found during the initial check, the following message is displayed:

- If you are using Windows, the following message is displayed:

  ```
  LaunchAnywhere Error: Could not find a valid Java virtual machine to load.

  You may need to reinstall a supported Java virtual machine.
  ```

- If you are using UNIX or Linux, the following message is displayed:

  ```
  No Java virtual machine could be found from your PATH
  environment variable.  You must install a VM prior to
  running this program.
  ```

After you ensure that Java 8 or later is installed, complete one of the following actions to correct the `Java Virtual Machine Not Found` error:

- Run the DS CLI installer again from the console, and provide the path to the JVM using the LAX_VM option. The following examples represent paths to the correct version of Java:
  - For a Windows system, specify the following path:

    ```
    dsclisetup.exe LAX_VM "C:\Program Files\java-whatever\jre\bin\java.exe"
    ```

    Note: Due to a space in the Program Files directory name, you are required to add quotation marks around the directory name.
  - For a UNIX or Linux system, specify the following path:

    ```
    dsclisetup.bin LAX_VM /opt/ibm-Java-whatever/java/bin/java
    ```

Note: If you use the LAX_VM argument, the installer attempts to use whatever JVM that you specify, even if it is an unsupported version. If an unsupported version is specified, the installation might complete successfully, but the DS CLI might not run and return an `Unsupported Class Version Error` message. You must ensure that you specify a supported version.

– Continue with the installation of the DS CLI.

- (For UNIX or Linux) Add the Java virtual machine location to your PATH environment variable by running the following command:

  ```
  export PATH=$PATH:/opt/ibm-Java-whatever/java/bin
  ```

  Then, run the dsclisetup.bin program to install the DS CLI.
- (AIX only) Run the following commands to sequentially disable the LIBPATH environment variable, install the DS CLI, and restore the LIBPATH environment variable:

  ```
  export LIBSAVE=$LIBPATH
        unset LIBPATH
  dsclisetup.bin LAX_VM/opt/ibm-Java-whatever/java/bin/java
        export LIBPATH=$LIBSAVE
        unset LIBSAVE
  ```

Upgrading the DS CLI on your system

You can upgrade the DS CLI on your system by following the removal and installation procedures.

There are considerations and preparation that you must make before you make this upgrade.

Notes:

1. The DS CLI installation program changed from the InstallShield MultiPlatform installer to the InstallAnywhere installer, which reduced many restrictions.
2. Installations of multiple DS CLIs using the InstallShield MultiPlatform installer are not supported.
3. Installations of multiple DS CLIs using the InstallAnywhere installer (with up to one DS CLI using the ISMP installer is supported).
4. As a result of supporting multiple DS CLI installations, the DS CLI location is no longer added to the Windows PATH environment variable. This might cause any existing shell scripts to fail if they rely on the DS CLI program location being in the PATH variable. To fix this problem, you can choose from one of the following solutions:
   - You can set the default DS CLI location by manually adding it to the PATH environment variable.
   - You can use the full DS CLI path name in all of the DS CLI scripts that you use.
   - You can write a script to change to a specific DS CLI location, then run the rest of your scripts by using that DS CLI location.
   - You can change the PATH environment variable so it is using the correct DS CLI location before calling each script.

As part of your upgrade preparation, consider the following items:

- If the DS CLI was previously installed with the InstallShield MultiPlatform installer, an upgrade requires that you remove the existing DS CLI and that you install the upgraded DS CLI. This method is the most certain way to ensure that you receive an error-free installation. However, this removal and installation process can be a concern where you have customized the system profile file. A reinstallation can most likely overwrite your current system profile file. If you want to keep your current system profile file, complete the following tasks:
  1. Make a copy of your current system profile file and save it in a convenient place.
  2. Merge the saved system profile file into the new system profile file in the DS CLI installation directory after the installation completes. Therefore, you can keep any customized variables, and retain any new variables in the system profile file that was installed with the upgraded DS CLI.

Note: Personal profiles that are not saved under the DS CLI installation directory are not affected by the upgrade process.

Uninstalling the DS CLI

You can uninstall the DS CLI by using the same modes that are allowed by the operating systems during the installation process. For example, you can use the graphical (swing) mode, unattended (silent) mode, or console mode to install this interface. Conversely, you can remove this interface by using the graphical (swing) mode, unattended (silent) mode, or console mode.

The following topics describe the steps that are required to successfully remove the DS command-line interface.

Overview of uninstallation methods

The following information provides an overview of the uninstallation methods that are supported and how to access them.

Uninstalling the DS CLI by using the graphical mode

Complete the following steps to uninstall the DS CLI from your system when the DS CLI is installed on a Windows or UNIX system.

Note: If you do not want to create a new profile when you reinstall the CLI, select to not delete the DS CLI profile as you complete this task, or copy the profile file to a safe place before you uninstall the CLI.

1. Navigate to the location where the DS CLI was installed. For example, on Windows the path might be C:\Program Files\ibm\dscli. On UNIX or Linux systems, the location might be /opt/ibm/dscli.
2. Open the `_uninst` directory to access the uninstaller. (On Windows, C:\Program Files\ibm\dscli_uninst\uninstaller.exe; on UNIX or Linux systems, /opt/ibm/dscli/_uninst/uninstaller.)
3. From the command prompt, specify the -i swing parameter to open the uninstaller in graphical mode. The Uninstall Storage Command Line Interface window is displayed. Click **Uninstall** to complete the uninstallation process, or Cancel to cancel the uninstallation.

Alternately, you can use the Add/Remove Programs facility of the Windows operating system to uninstall the DS CLI from your system.

Uninstalling the DS CLI by using unattended (silent) mode

Use the unattended (silent) mode to uninstall the DS CLI through the command line if the DS CLI is installed on any system other than IBM i i5/OS.

Open a command prompt and type: **<install_directory_path>/_uninst/uninstaller -i silent**. Then press **Enter** to remove the DS CLI.

Uninstalling the DS CLI by using the console mode

Use the console mode to uninstall the DS CLI when the DS CLI is installed on a UNIX system that does not have an X display.

Open a command prompt and type: **<install_directory_path>/_uninst/uninstaller -i console**. Then press **Enter** to display the Uninstall Storage Command Line Interface page. Press **Enter** again to complete the uninstallation.

Uninstalling the DS CLI from an IBM i model

Because the DS CLI is installed on an IBM i model from a remote system, it is not possible to use the conventional DS CLI removal methods that you use with other systems.

When the DS CLI was installed onto your IBM i model, you used a remote system for the installation (for example, Windows, UNIX or AIX). Part of the installation process is the creation of an uninstaller. However, because you were using another system to do your installation, the uninstaller that was created was for the system that you installed from and not for the IBM i model. This uninstaller cannot be used to uninstall the DS CLI.

When you want to uninstall the DS CLI, you can use one of the following two methods:

- Uninstall directly from your i5/OS IBM i model by completing the following steps:
  1. Delete the library by using DLTLIB QDSCLI.
  2. Run the command, EDTF 'DSCLI_INSTALL_PARENT', where 'DSCLI_INSTALL_PARENT' is the parent directory of the DS CLI installation. The default parent directory is /ibm.
  3. Insert a 9 (recursive delete) beside the DS CLI directory to remove all DS CLI java code.

  You might use this method if you are not planning to upgrade the DS CLI and you want to totally uninstall the DS CLI from your IBM i model.
- Uninstall by using a remote system.

  You might use this method when you are upgrading the DS CLI, because after the removal, you can use this remote system to install the upgraded DS CLI.

Uninstalling the DS CLI using your IBM i model directly

Note: The i5/OS direct removal method requires that you use the i5/OS console mode and that you use an i5/OS command. The following steps presume that you are logged in to the i5/OS and have the authority to remove applications.

To uninstall the DS CLI using your IBM i model directly, enter the following command from your i5/OS application:

```
RUNJVA CLASS(run) PARM('-console')
CLASSPATH('/Q IBM/ProdData/Java400/jt400ntv.jar:/yourdir/_uninst
/uninstall.jar')
```

Substitute your uninstall directory for *yourdir*.

Uninstalling the DS CLI from your IBM i model using the remote method

To remove the DS CLI from your IBM i model using the remote method, go to the `_uninst` folder on the Windows desktop and run the uninstaller.exe file.

Running and configuring the DS CLI

Complete these tasks to set up the DS CLI so that you can use the DS CLI to configure your DS8000.

Before you begin

You must install the DS CLI before you complete these postinstallation tasks. See Chapter 2, “Running the DS CLI,” on page 25 for information about running the DS CLI, obtaining help, and interpreting exit codes.

About this task

Note: Before you can use the DS CLI with the i5/OS, if you are using external load source, you must set up the initial configuration of your DS8000. After the initial configuration, you can do a D-mode IPL and begin using the DS CLI directly from the i5/OS. In the meantime, you can follow the list below as a guide for your initial configuration.

Procedure

Complete these tasks to complete the installation of the DS CLI:

1. Set your DS CLI default configuration settings.
   - If this is a new installation, complete the rest of this procedure.
   - If this is not a new installation, do not complete the following steps.
2. Initiate the DS CLI to begin using it in either single-shot, script, or interactive command mode.
3. Set up your required user accounts and passwords.
4. Activate your licensed functions. This includes obtaining your feature activation codes and applying the feature activation codes to your storage unit.
5. Use the DS CLI to enable SNMP traps for Copy Services events and storage complex events on your storage unit.
6. Optionally, configure encryption on your encryption-capable storage unit.
7. Configure new fixed block or CKD storage. Use the DS CLI to create and modify fixed block extent pools, arrays, ranks, volumes, and volume groups. You can also configure host ports and connections.

Creating a default CLI profile

You can specify default settings for the CLI by defining one or more profiles on the system. For example, you can specify the default output format for list commands, the primary and secondary storage manager IP addresses for the DS8000, or the storage image ID that is required by many commands.

If a user enters a value with a command that is different from a value in the profile, the command overrides the profile.

The following options are available for profile files:

- You can modify the default profile. The default profile, dscli.profile, is installed in the profile directory with the software. For example, `c:\Program Files (x86)\IBM\dscli` is the directory path for operating systems Windows 7 and later. The directory path for operating systems UNIX and Linux is `/opt/ibm/dscli/profile/dscli.profile`.

  Note: Changing the default profile changes the DS CLI default settings for all users. If you do not want to change the DS CLI default settings for all users, consider creating a personal default profile instead.
- You can create a personal default profile by making a copy of the system default profile as <*user_home*>/dscli/profile/dscli.profile.
- You can create a profile for a specific DS8000 system by making a copy of the system default profile and specifying the primary and secondary management console IP addresses and the storage image ID. For example:

  <*user_home*>\dscli\profile\DS8000_name1

  <*user_home*>\dscli\profile\DS8000_name2
- You can create a profile for the storage unit operations, typically for Copy Services commands, by starting with a specific DS8000 profile and then adding the remote storage image ID. For example:

  <*user_home*>\dscli\profile\operation_name1

  <*user_home*>\dscli\profile\operation_name2

These profile files can be specified using the DS CLI command parameter **-cfg <*profile_name*>**. Profile names are not required to use the .profile file name extension or any extension. However, the -cfg *profile_name* parameter must be a complete file name, including the extension if one is specified. Also, if the profile is stored in the user's personal profile directory at <*user_home*>\dscli\profile, you need to specify only the file name. If the profile is stored in any other directory, the <*profile_name*> must also include the full path name. If the -cfg profile file is not specified, the user's default profile file is used. If a user's profile file does not exist, the system default profile file is used.

The home directory <user_home> is defined by the Java system property named user.home. The location of your home directory is determined by your operating system. The following examples are home directories in different operating systems:

Windows 7 and later operating systems
For Windows 7 and later operating systems, the property value defaults to the environment variable **%USERPROFILE%** in a directory called c:\Users\Administrator.

UNIX or Linux operating systems
For a UNIX or Linux operating system, the property value defaults to the environment variable **$HOME**. As a result, your personal profile is ~/dscli/profile/dscli.profile.

When you install the DS CLI, the default profile is installed in the profile directory with the software. The file name is dscli.profile; for example, c:\Program Files (x86) \IBM\dscli\profile\dscli.profile.

The following steps provide a Windows example of the process you can use to modify key items in the profile file:

1. Click the DSCLI icon on your desktop. A command prompt window is displayed.
2. Enter cd profile at the command prompt to move to the system default profile directory.
3. Edit the profile file with a text editor such as NotePad or WordPad.
4. Scroll down to the number (#) sign in front of hmc1: and remove the # sign.
5. Enter the correct IP address of your management console.
6. If this is a dual HMC DS8000, perform steps 4 and 5 for hmc2.
7. Scroll down to the # sign in front of devid and remove the # sign.
8. Enter the serial number of your machine type (include the values for manufacture, machine type, and serial number). For example, IBM.2107-75YZ881.
9. Save the file.
10. Enter cd.. at your command prompt.
11. Enter DSCLI at your command prompt and the DS CLI starts. You are asked to provide only your user ID and password and not the address of your management consoles.

Table 6 provides the list of profile variables that you can use to create the profile.

Table 6. Profile variables

Variable	Description
username: *string*	Specifies your user name for entering DS CLI commands. This variable is equivalent to option **-user**.
password: *string*	Specifies the password to authenticate when you start a CLI session. This parameter is not required nor recommended. If you use this method to designate your password, the password is displayed on the screen. Another option is to specify a password file (pwfile) that is used when you start the DS CLI. This variable is equivalent to option **-passwd**.
pwfile: *passwordfile*	Specifies a password file containing your password as an alternative to the variable *password*. This variable is equivalent to option **-pwfile**.

Table 6. Profile variables (continued)

Variable	Description
banner: on | off	Enables or disables the banner that appears before the command output. This variable is equivalent to the command option **-bnr**. The command option **-bnr** overrides this default value.
delim: *character*	Specifies a delimiter character for the format: delim variable. The default character is a comma. This variable is equivalent to the command option **-delim**. The command option **-delim** overrides this default value.
devid: *string*	Specifies the storage image ID that is the target for the command. This value is equivalent to the command option **-dev**. The command option **-dev** overrides this default value.
echo: on | off	Specifies whether the command is printed before it is executed. Specify one of the following formats: • on: Specifies that the command is printed before it is executed. • off: Specifies that the command is not printed before it is executed.
echoprefix: *prefix* | none	Specifies the command prefix to print before a command is executed. • echoprefix: Specifies the prefix to print before a command is executed. If **echo** is on and **echoprefix** is specified, then its value is to be printed on the line before the echoed command. • none: Specifies that no prefix is to be printed before an echoed command.
format: *option*	Specifies the output format for list commands. Specify one of the following formats: • default: Specifies default output. • xml: Specifies XML format. • delim: Specifies columnar format. Columns are delimited with the character that you must specify with the delim variable. • stanza: Specifies a vertical table. This variable is equivalent to the command option **-fmt**. The command option **-fmt** overrides this default value.
fullid: on | off	Specifies that IDs display in fully qualified format, which includes the storage image ID.
header: on | off	Enables or disables the headers that display with the columns of data in the list commands. This variable is equivalent to the command option **-hdr**. The command option **-hdr** overrides this default value.
hmc1: *string*	Specifies the primary Storage Manager IP address. This variable is equivalent to the command option **-hmc1**. The command option **-hmc1** overrides this default value.
hmc2: *string*	Specifies the secondary Storage Manager IP address. This variable is equivalent to the command option **-hmc2**. The command option **-hmc2** overrides this default value.

Table 6. Profile variables (continued)

Variable	Description
locale: *code*	Specifies the language for the output on the local computer. • ar: Arabic • be: Byelorussian • bg: Bulgarian • ca: Catalan • cs: Czech • da: Danish • de: German • el: Greek • en: English • es: Spanish • et: Estonian • fi: Finnish • fr: French • gu: Gujarati • hi: Hindi • hr: Croatian • hu: Hungarian • in: Indonesian • is: Icelandic • it: Italian • iw: Hebrew • ja: Japanese • kk: Kazakh • kn: Kannada • ko: Korean • lt: Lithuanian • lv: Latvian (Lettish) • mk: Macedonian • mr: Marathi • ms: Malay

Table 6. Profile variables (continued)

Variable	Description
locale: *code*	• nl: Dutch • no: Norwegian • pa: Punjabi • pl: Polish • pt: Portuguese • ro: Romanian • ru: Russian • sa: Sanskrit • sh: Serbo-Croatian • sk: Slovak • sl: Slovenian • sq: Albanian • sr: Serbian • sv: Swedish • ta: Tamil • te: Telugu • th: Thai • tr: Turkish • uk: Ukrainian • vi: Vietnamese • zh: Chinese
maxNumReports: *number*	Sets the maximum number of records (lines) for an I/O Performance Manager performance report. **Note:** The default maximum number of records for a performance report is 256. The value for maxNumReports is recommended to be no larger than 3000. If the target is a DA pair, the recommended value is to be no larger than 1500.
port: 1718 | 1750 | 1751	Specifies the port that the DS CLI should use when connecting to the DS8000 system. If the port is not specified, the DS CLI first attempts to connect using port 1751 with a NIST-compliant certificate. If that connection attempt fails, it attempts to connect to the existing DS8000 port 1750 with the legacy certificate. If the second attempt fails, the DS CLI attempts to connect to port 1718 with the legacy certificate used by ESS 2105 machines. This default behavior means that the DS CLI will connect to any ESS 2105 or DS8000 system. This default behavior means that the DS CLI will connect to any ESS 2105 or DS8000 system. However, checking multiple ports can cause a connection delay when a Release 7.2 or later DS CLI attempts to connect to a DS8000 system or ESS 2105 machine that does not listen on the 1751 port. To prevent the additional delay, you can use this variable to specify a single attempt on the specified port. **1718** Attempt to connect using only port 1718 (ESS 2105 with legacy certificate). **1750** Attempt to connect using only port 1750 (DS8000 prior to Release 7.2 with legacy certificate). **1751** Attempt to connect using only port 1751 (DS8000 Release 7.2 and later with NIST SP 800-131a-compliant certificate).
paging: on | off	Controls the display of output. If paging is enabled, a limited number of lines of output displays when a command is issued. The lines do not scroll. You must set the number of lines per page with the rows variable. This variable is equivalent to command option **-p**. The command option **-p** overrides this default value.

Table 6. Profile variables (continued)

Variable	Description
timeout: *number*	Sets the timeout value of client/server synchronous communication. The unit of the value is seconds. The default value is 900 seconds. You can set this timeout if the processing of a command ends by timeout due to network or client or server performance issue. **Note:** The command timeout value can be longer than this value because one command can consist of multiple client/server requests.
timeout.connection: *number*	Sets the timeout value to establish client or server connection. The unit of this value is seconds. The timeout value must be greater than zero. System-default socket timeout value is used if the value is set to zero. The default value is 20 seconds. **Notes:** 1. If the DS CLI returns a connection error, check for the following conditions: • Is there a secure physical connection between the client and server? • Is the default timeout value too short to establish a connection? 2. Setting a connection timeout value that is too short can cause unexpected connection problems.
remotedevid: *string*	Specifies the remote storage image ID. This variable is equivalent to the command option **-remotedev**. The command option **-remotedev** overrides this default value.
rows: *number*	Specifies the number of rows per page of output if the paging variable is enabled. This variable is equivalent to command option **-r**. The command option **-r** overrides this default value.
verbose: on | off	Enables or disables verbose output. This variable is equivalent to the command option **-v**. The command option **-v** overrides this default value.

Example

```
#
# DS CLI Profile
#

#
# Management Console/Node IP Address(es)
#   hmc1 and hmc2 are equivalent to -hmc1 and -hmc2 command options.
#hmc1: 127.0.0.1
#hmc2: 127.0.0.1

#
# Default target Storage Image ID
#    "devid" and "remotedevid" are equivalent to
#    "-dev storage_image_ID" and "-remotedev storage_image_ID" command options,
#     respectively.
#devid:  IBM.2107-AZ12341
#remotedevid: IBM.2107-AZ12341

#  pwfile
#  Specifies a password file containing your password as an alternative
#  to the variable of password.
#  pwfile is equivalent to command option -pwfile
#  Example: pwfile:c:/mydir/75CNF11/pwfile.txt
#
# locale
#    Default locale is based on user environment.
#locale:  en

# Timeout value of client/server synchronous communication in second.
```

```
#    DSCLI command timeout value may be longer than client/server communication
#    timeout value since multiple requests may be made by one DSCLI command
#    The number of the requests made to server depends on DSCLI commands.
#    The default timeout value is 900 seconds.
#timeout:  900

# Socket connection timeout value in seconds.
#    The timeout value must be greater than zero.
#    System default socket timeout value is used if timeout value is set to zero.
#    The default connection timeout value is 20 seconds.
#timeout.connection: 20

# Output settings
#
# ID format of objects:
#   on:  fully qualified format
#   off: short format
fullid:  off

# Paging and Rows per page.
# paging enables/disables paging the output per line numbers specified by "rows".
# "paging" is equivalent to "-p on|off" option.
#   on  : Stop scrolling per output lines defined by "rows".
#   off : No paging. (default)
# "rows" is equivalent to "-r #" option.
paging:  off
#rows:  24

# Output format type for ls commands, which can take one of the following values:
#   default: Default output
#   xml    : XML format
#   delim  : delimit columns using a character specified by "delim"
#   stanza  : Horizontal table format
# "format" is equivalent to option "-fmt default|xml|delim|stanza".
#format:  default

# delimiter character for ls commands.
#delim:  |
# Display banner message. "banner" is equivalent to option "-bnr on|off".
#   on  : Banner messages are displayed. (default)
#   off : No Banner messages are displayed.
banner:  on

#
# Display table header for ls commands. "header" is equivalent
# to option "-hdr on|off".
#   on  : Table headers are displayed. (default)
#   off : No table headers are displayed.
header:  on

#
# Display verbose information. "verbose" is equivalent to option "-v on|off".
#   on  : Display verbose information.
#   off : No verbose information.
verbose: off

# Echo each dscli command.
# on : Echo commands to standard out prior to execution. Passwords within
command line arguments will be hidden.
# off : No command echo. (default)
#echo:on

# If echo is on and echoprefix is specified, its value will be printed on the
```

```
line before the echoed command.
#echoprefix:dscli>

# The max number of records for performance report.
# The default max number of records for performance report is 256.
# The value for it is suggested to be
# not larger than 3000. If the target is dapair, the value is
# suggested to be not larger than 1500.
#maxNumReports: 256

# Connection port number used when connecting to the DS8000.
# This is equivalent to –port 1718 | 1750 | 1751
# on the command line. If not specified,
# the DSCLI first attempts to connect using the new port 1751
# with a NIST-compliant certificate, and if that fails, it attempts to connect to existing
# DS8000 port 1750 with the legacy certificate. If the second attempt also fails, the DSCLI
# attempts to connect to port
# 1718 with the legacy certificate used by ESS 2105 machines. While this default
# behavior means that the R7.2+ DSCLI will connect to any ESS 2105 or DS8000, checking multiple
# ports can cause a connection delay when a R7.2+ DSCLI attempts to connect to a DS8000 or ESS
# 2105 that does not listen on the 1751 port. To prevent this additional delay, this variable
# may be used to specify a single attempt on the specified port.
# 1718 : Only attempt to connect using port 1718 (ESS 2105 with legacy certificate).
# 1750 : Only attempt to connect using port 1750 (DS8000 prior to R7.2 with legacy certificate).
# 1751 : Only attempt to connect using port 1751 (DS8000 R7.2+ with NIST compliant certificate).
#port: 1750

# End of Profile
```

Setting up user accounts by using the DS CLI

This task describes how to set up a user account. You must have administrator authority to enable this function.

Before you begin

The default administrator and security administrator accounts are set up automatically at the time of installation. To access the storage administrator account, use the user name *admin* and the default password *admin*. To access the security administrator account, use the user name *secadmin* and the default password *secadmin*. These passwords are temporary and expire after their initial use. You must change the password before you can use any of the other functions. The storage administrator can assign a user to one or more user roles, except for the security administrator role. Only the security administrator can assign a user to the security administrator role. The user roles and the associated functions that are allowed by the assignment are as follows:

admin (Administrator)
: All users that you assign to the storage administrator user role have access to all DS8000 storage image resources except those that are reserved for security administrator users.

ibm_engineering) (Engineering)
: This user role is typically assigned to IBMsupport personnel that perform all service functions and other functions that might be needed. This role does not have access to the logical configuration or data on the storage system.

ibm_service) (Service)
: This user role is typically assigned to IBMsupport personnel that service the hardware (install, remote, or repair) and update firmware. This role does not have access to the logical configuration or data on the storage system. Users in this role cannot be assigned to any other role, and users in any other role cannot be assigned to this role.

op_volume (Logical Operator)
The logical operator user role allows access to service methods and resources that relate to logical volumes, hosts, host ports, logical subsystems, logical volumes, and volume groups, excluding security methods. In addition, this user role inherits all authority of the monitor user role.

op_storage (Physical Operator)
The physical operator user role allows access to physical configuration service methods and resources, including storage complex, storage image, array, rank, and extent pool objects. This user role inherits all the authority of the Copy Services operator and logical operator user roles, excluding security methods.

op_copy_services (Copy Services Operator)
The Copy Services operator user role allows access to all Copy Services service methods and resources, excluding security methods. In addition, this user role inherits all authority of the monitor user role.

monitor (Monitor)
The monitor user role allows access to list and show commands. It provides access to all read-only, nonsecurity management console-server service methods and resources.

no_access (No Access)
The no_access user role does not allow access to any service methods or storage image resources. By default, this user role is assigned to any user account in the security repository that is not associated with any other user role.

secadmin (Security Administrator)
All users that you assign to the security administrator user role can initiate recovery key operations, and add other users to this role. Users in this role can not be assigned to any other user role, and users in any other user role can not be assigned to this role.

Table 7. Permissions for roles

Permission	Admin-istrator	Copy operator	IBM Engineering	IBM Service	Logical operator	Logical and copy operator	Physical operator	Security admin-istrator
Arrays and Pools								
Create pools and assign arrays	X						X	
Manage pools and arrays	X						X	
Delete pools and unassign arrays	X						X	
IBM Z® Volumes and LSSs								
Configure IBM Z volumes and LSSs	X				X	X	X	
Manage IBM Z volumes and LSSs	X				X	X	X	
Delete and reinitialize IBM Z volumes	X				X	X	X	
Open System Volumes								

Table 7. Permissions for roles (continued)

Permission	Administrator	Copy operator	IBM Engineering	IBM Service	Logical operator	Logical and copy operator	Physical operator	Security administrator
Configure open system volumes	X				X	X	X	
Manage open system configuration	X				X	X	X	
Delete open system configuration	X				X	X	X	
IBM i Volumes and LSSs								
Configure IBM i configuration	X				X	X	X	
Manage IBM i configuration	X				X	X	X	
Delete and reinitialize IBM i configuration	X				X	X	X	
Local and Remote Access Management								
Manage security administrators								X
Manage local user accounts	X							
Manage user roles	X							
Manage remote authentication	X							
Manage network security	X		X	X			X	
Notifications and Monitoring								
Manage SNMP trap notifications	X							
Manage system events	X							
Offload audit log	X							
Manage call home	X							
Manage syslog settings	X		X	X				

Table 7. Permissions for roles (continued)

Permission	Admin-istrator	Copy operator	IBM Engineering	IBM Service	Logical operator	Logical and copy operator	Physical operator	Security admin-istrator
Remote Support								
Manage Remote Support Center (RSC) connection	X							
Manage Assist On-Site (AoS) connection	X		X	X				
Change HMC access settings	X		X	X			X	
Trouble-shooting	X		X	X				
System Settings								
Manage Fibre Channel port settings	X		X	X			X	
Manage Fibre Channel port security	X							
Manage system settings	X		X	X			X	
Install software	X							
Power off system	X		X	X	X	X	X	
Encryption								
Manage data at rest encryption	X		X					
Create data at rest recovery key								X
Manage Fibre Channel Port Endpoint Security	X		X					
Feature Settings								
Modify Ethernet settings	X		X	X				
Modify Easy Tier settings	X		X				X	

Table 7. Permissions for roles (continued)

Permission	Admin-istrator	Copy operator	IBM Engineering	IBM Service	Logical operator	Logical and copy operator	Physical operator	Security admin-istrator
Modify zHyperLink settings **Note:** Not available on DS8882F systems.	X							
Manage resource group settings	X		X	X			X	
Manage performance group settings	X							
Manage cloud settings	X							
FlashCopy								
Manage FlashCopy relationships	X	X				X		
Manage remote FlashCopy relationships	X	X				X		
Mirroring and Paths						X		
Manage mirroring paths	X	X				X		
Manage mirroring relationships	X	X				X		
Modify LSS CS settings	X	X				X		

In addition to assigning users to one or more user roles, you also must assign a default password to each user. When you notify users of their role assignment and default password, indicate that the default password is only good for the initial logon. Users must change the password at the time of their initial logon. Also, remind all users to record their password in a safe place because there is no way that the administrator or the application can retrieve a password.

Note: You must change the default password for an account, including the administrator account, to be able to use any CLI command other than the one to change the password. See the **chuser** command for more information.

About this task

Use the **mkuser** DS CLI command to create new user accounts with specific roles (user role or roles) and an initial password. If you assign multiple roles to an account, ensure that you separate the different roles with a comma. For example, `op_volume, op_storage`. See the **mkuser** command description for more details.

Procedure

1. Log in to the DS CLI in interactive command mode.
2. Type the following command from the dscli command prompt to assign a user to an account with a default password: dscli> `mkuser -pw AB9cdefg -group service,op_copy_services -pol my_policy1 testuser`
3. Press **Enter** and observe the processing result. A successful process returns the following display:
   `User Name testuser with my_policy1 successfully created.`

Activating your machine and feature licenses by using the DS CLI

Use the steps that are described in this task to activate your license activation codes. These codes must be activated before any configuration can be applied to your DS8000 network.

Before you begin

See the *IBM DS8880 Introduction and Planning Guide* (GC27-8525) for information.

Multiple codes are associated with these features. For a complete list of activation codes, see the IBM Disk Storage Feature Activation (DSFA) website at: Data storage feature activation (www.ibm.com/storage/dsfa/)

On the website, enter your machine type, serial number, and machine signature. If you need to acquire the information to complete these fields, issue the DS CLI **showsu** command to show the machine type and serial number, and the **showsi** command to show the machine signature.

After the activation codes are available, download the codes in XML format onto a CD or USB drive. You can then import the codes from the XML file when you process the DS CLI **applykey** command.

Notes:

1. For DS8000, in most situations, the DSFA application can locate your 283x license authorization record when you enter the DS8000 (2107) serial number and signature. However, if the 2244 license authorization record is not attached to the 2107 record, you must assign it to the 2107 record in the DSFA application. In this situation, you must have the 283x serial number (which you can find on the License Function Authorization document).

About this task

The DS CLI **applykey** command activates the licenses for your storage unit. The DS CLI **lskey** command verifies which type of licensed features are activated for your storage unit.

Complete the following steps to activate your license activation codes:

Procedure

1. Log in to the DS CLI in interactive command mode.
2. Issue the DS CLI **applykey** command at the dscli command prompt as follows. (This example presumes that your XML file is named "keys.xml" and it resides on a CD or USB drive): `dscli> applykey -file a:\keys.xml -dev IBM.2107-75FA120`
3. Press **Enter**. When the process completes, the following message is displayed:
   `Licensed Machine Code key xxxx, key xxxx successfully applied.`
4. Verify that the keys are activated for your storage unit by issuing the DS CLI **lskey** command as follows: `lskey -dev IBM.2107-75FA120`
5. Press **Enter** and the following type of report is displayed:

Activation key	Authorization level (TB)	Scope
Operating environment (OEL)	45	All
Remote mirror and copy (RMC)	25	All
Metro mirror (MM)	25	All
Global mirror (GM)	25	All
Metro/global mirror (MGM)	25	All
Remote mirror for z/OS (RMZ)	25	CKD
Point in time copy (PTC)	25	All
Parallel access volumes (PAV)	100	CKD
IBM HyperPAV	On	CKD
IBM FlashCopy SE	105	All

Chapter 2. Running the DS CLI

You can use the DS CLI in three different command modes.

The following command modes are available:
- Single-shot
- Interactive
- Script

Note: DS CLI logs are created on the system running the DS CLI. The location of the logs is under the user home directory at `dscli/log/`. This location cannot be changed. For example, on Windows the logs for the administrator user are created under the directory `C:\Documents and Settings\Administrator\dscli\log`. On UNIX, the logs are created under the directory `~/dscli/log`.

Logging in to the DS CLI

You must log in to the DS CLI to use any of the command modes.

Before you begin

You must ensure that you are in the directory where you installed the DS CLI, or use the full path name of the DS CLI. The following list provides the location of the default directory for each operating system:

AIX /opt/ibm/dscli

Linux /opt/ibm/dscli

Oracle Solaris
/opt/ibm/dscli

Windows
(32-bit) C:\Program Files\IBM\dscli

(64-bit) C:\Program Files (x86)\IBM\dscli

IBM i /ibm/dscli

z/OS /opt/IBM/CSMDSCLI/dscli

About this task

To log in to the DS CLI, open a command prompt and enter `dscli` with the following information:

HMC1
Specify the primary management console.

Note: If you are using a 2105 machine type as part of your network and are going to use the Copy Services functions, you must specify the IP address of the primary or secondary domain control server where you installed the DS CLI.

User Name
Specify the name of the user account. The default user names for the first logins are *admin* and *secadmin*.

© Copyright IBM Corp. 2004, 2019

Password

Specify the user password. The default password for the administrator account is *admin* and the default password for the security administrator account is *secadmin*. However, these passwords are only good for the first login.

Note: Because the passwords for the administrator and security administrator accounts expire after you log in for the first time, you must change the password before you can perform any other DS CLI command function. Use the **chuser** command to change your password.

Each time you log in to the DS CLI, you can specify this information by using either of the following three methods:

- You can log in to the DS CLI without specifying any of this information on the command line and the application prompts you to enter the information interactively. For example: `/opt/ibm/dscli/dscli`.
- You can log in to the DS CLI by specifying this information on the command line. For example: `dscli -hmc1 mtc032h.storage.tucson.nnn.com -user admin -passwd topn0t`.
- You can log in to the DS CLI by specifying this information on the command line, except for the password, by using the **-pwfile** parameter instead of the **-passwd** parameter. For example: `dscli -hmc1 9.1.12.123 -user admin -pwfile /home/ming/dscli/security.dat`, where "ming" is the user ID used to log in to the operating system. The security.dat file was created by using the DS CLI command **managepwfile**. See the DS CLI command **managepwfile** for more details.

Notes:

1. Entering a DS CLI command at the **dscli** command prompt requires that you continue entering all the parameters and values until the command is complete. This can result in an automatic line wrap if your command has many parameters and values.
2. You cannot force a line break or wrap by pressing the Enter key and then typing the rest of the command on a second line. The DS CLI interprets the use of the Enter key as an end to the function and begins to process whatever is contained on the line, ignoring the second line.
3. The DS CLI command examples are often shown with line wraps that would not occur during your input. These examples are displayed for clarity and other formatting considerations.

Using the DS CLI single-shot command mode

Use the DS CLI single-shot command mode if you want to enter an occasional command but do not want to keep a history of the commands that you entered previously.

About this task

You must supply the login information and enter the command that you want to process at the same time.

Complete the following steps to use the single-shot mode:

Procedure

1. Use the following command format to start a DS CLI session (Windows operating system):

   ```
   dscli -hmc1 192.168.30.15 -user admin -passwd my_password lssi -s -fullid
   -hdr off
   ```

   An example of this same command in IBM i without the report delimiters:

   ```
   DSCLI SCRIPT(*NONE) HMC1('192.168.30.15') USER(admin) PASSWORD(my_password)
   DSCLI(lssi)
   ```

   This command demonstrates the use of the **lssi** command with the **-s** parameter. Use this command to view the storage image IDs for your storage complex. The storage image ID consists of the manufacturer name (IBM), the machine type (2107 or 1750) and the serial number.

Notes:

a. The command example uses the **-fullid** DS CLI command parameter. The **-fullid** command parameter generates fully qualified IDs, which include the storage image ID, for every ID that is displayed in the command output.

b. The command example also uses the **-hdr** *off* command parameter, which turns off the header that is associated with the report that is generated from the **lssi** command.

c. Almost every DS CLI command requires the use of the storage image ID. You can set it as an environment variable by setting *devid* in the profile file, or by setting *devid* by using the **setenv** command. If you choose not to set it, a default *devid* is used on systems where the management console is aware of only one storage image. If you provide the **-dev** (*storage_image_ID*) parameter in commands, the value that you type takes priority over the *devid* environment variable. If you specify a full ID that contains the storage image ID, the storage image ID that you specify takes priority over the value from the **-dev** (*storage_image_ID*) parameter and over the *devid* environment variable.

2. Wait for the command to process. The following type of report is generated to list the storage image IDs that are associated with the storage complex.
   - IBM.2107-75FA111
   - IBM.2107-75FA112

Using the DS CLI script command mode

Use the DS CLI script command mode if you want to enter a sequence of DS CLI commands. Administrators can use this mode to create automated processes. For example, this mode can be used to establish remote mirror and copy relationships for volume pairs.

Before you begin

Consider the following requirements when you use the DS CLI script command mode:

- The DS CLI script can contain only DS CLI commands. Use of shell commands results in a process failure.
- You can add comments to the scripts. Comments must be prefixed by the number sign (#). For example, `# This script contains PPRC Path establish procedures`.
- Set the *echo* environment variable by using the dscli profile file or use the **setenv** command to specify that the dscli command name is displayed before the command output.

Note: An example script is displayed for your use as a guide.

About this task

You can enter the DS CLI script from the command prompt at the same time that you provide your login information.

Procedure

1. Enter the script name at the command prompt by using the following format:

   ```
   dscli -hmc1 9.1.23.45 -user admin -passwd password
   -script ~/bin/mkpprcpairs
   ```

   Note: If you are using IBM i and logged on to the DS CLI, you start the script mode by using the following format:

   ```
   DSCLI SCRIPT('/myscript') USER(admin) OUTPUT('/outfile')
   ```

2. Wait for the script to provide success or failure of the process.

Example

The following example shows a script that is used to establish remote mirror and copy relationships for volume pairs for the DS8000 series:

```
mkpprc -dev IBM.2107-1303561 -remotedev IBM.2107-7504491 -type mmir 1000-103F:2300-233F
mkpprc -dev IBM.2107-1303561 -remotedev IBM.2107-7504491 -type gcp 1100-113F:2340-237F
mkpprc -dev IBM.2107-1303561 -remotedev IBM.2107-7504491 -type mmir 1800-187F:2800-287F
mkpprc -dev IBM.2107-1303561 -remotedev IBM.2107-7504491 -type gcp 1200-127F:2500-257F
mkpprc -dev IBM.2107-1303561 -remotedev IBM.2107-7504491 -type mmir 1040-1054:2700-2714
mkpprc -dev IBM.2107-1303561 -remotedev IBM.2107-7504491 -type gcp 1055-107F:2400-242A
mkpprc -dev IBM.2107-1303561 -remotedev IBM.2107-7504491 -type mmir 1140-117F:2600-263F
```

Using the DS CLI interactive command mode (history and reports)

Use the DS CLI interactive command mode when you must process multiple transactions that cannot be incorporated into a script. The interactive command mode provides a history function that makes repeating or checking prior command usage easy to do.

About this task

In addition to being able to enter DS CLI commands at the DS CLI command prompt, a history function provides a view of the last DS CLI commands that you used. You can also repeat any of the last commands more quickly than having to type out the entire command. The example at the end of this process shows how the history function works.

The Windows platform allows the user to retrieve prior commands by using the up arrow, but this is a feature of the Windows platform and is not built into the DS CLI. While other platforms might not allow the up arrow feature, all of the supported platforms have the history function described here.

Note: Many of the DS CLI commands have a feature that allows you to specify a dash (-) for the parameter and the specified value is read from standard input. However, you cannot use the dash (-) while in the DS CLI interactive command mode.

Complete the following steps to use the DS CLI in the interactive command mode:

Procedure

1. Log on to the DS CLI at the directory where it is installed.
2. Provide the information that is requested by the information prompts. The information prompts might not appear if you have provided this information in your profile file. The command prompt switches to a dscli command prompt.
3. Begin using the DS CLI commands and parameters. You are not required to begin each command with **dscli** because this prefix is provided by the dscli command prompt.

   Tip: Issue the **setenv** command to control how the reports that are generated by the **ls** commands are displayed on your computer. You can set or display command output format options by using the **setenv** command. For example, you can specify that the reports be displayed in one of the following formats:

   delim Displays output in a table format and sets the column delimiter to a single character.

   xml Displays output in XML format.

   stanza Displays output in stanza (vertical table) format.

   See the **setenv** command for more details.

Example

To use the DS CLI history function that is associated with the interactive command mode, complete the following steps:

1. Issue an exclamation mark (!) to display CLI commands that you used in the current session. For example: dscli> ! results in a list of commands such as the following example:

   [4] lsarraysite -dev IBM.2107-1300771
   [3] lsarray -dev IBM.2107-1300771
   [2] lsextpool -dev IBM.2107-1300771
   [1] lsextpool -dev IBM.2107-1300771
2. Issue dscli> !1 to reissue the last command. Or, issue dscli> !3 to reissue command [3].

Obtaining the serial (storage image ID) number by using the DS CLI

Almost every DS CLI command requires the use of the storage image ID. If you add your target storage image ID into your profile file under the *devid* designation, you are not required to provide the storage image ID when you issue each command.

About this task

Use the **lssi** DS CLI command to list the storage image IDs that are associated with your storage complex. The storage image ID consists of the manufacture name (IBM), the machine type (2107 or 1750), and the serial number. You can record the target storage image ID in your profile file. The **lssi** command can save you input time when you must process many transactions that cannot be part of a script.

To obtain a list of the storage image ID numbers, complete the following steps:

Procedure

1. Log in to the DS CLI in interactive command mode.
2. Enter the following command format at the dscli command prompt to obtain the storage image IDs:
   ```
   dscli> lssi -s
   ```
3. Wait for the command to process. The following type of report is generated, which lists the storage image IDs that are associated with the storage complex.
   - IBM.2107-75FA111
   - IBM.2107-75FA112

DS CLI command help

You can view online help for each CLI command. To view the help, type the word `help` and the command name at the dscli command prompt.

The **help** command contains parameters that influence the type of help information that you can receive:

Table 8. Parameters for the help command

Command	Description
help	Displays a list of all the DS CLI commands that are available for use.
help -s	Displays a list of commands with brief descriptions.
help -l	Displays a list of commands with their associated syntax.
help -match <string>	Displays a list of commands that contain the specified string. This parameter cannot be used with any other parameters.

Table 8. Parameters for the help command (continued)

Command	Description
command_name -h command_name -help command_name -? help command_name	Displays the reference page (man page) for the command name.
help -s command_name	Displays the brief description for the command name.
help -l command_name	Displays the usage statement for the command name.

Notes:

- You cannot use the **-s** and **-l** parameters with the following help command parameters: **-h**, **-help**, and **-?**.
- Much of the information that is associated with the **help** command is displayed in list format. You can include the page (-p on) and row (-r *number*) controls; for example, `dscli> help -p on -r 20`. This command pauses your page listing after 20 entries and prompts you to `press any key to continue`.

Example

```
dscli> help -match flash
```

The resulting output

```
reverseflash
commitflash
rmflash
lsremoteflash
resyncremoteflash
resyncflash
mkremoteflash
setflashrevertible
revertflash
unfreezeflash
revertremoteflash
rmremoteflash
commitremoteflash
mkflash
lsflash
setremoteflashrevertible
```

Obtaining and interpreting DS CLI exit codes

Complete this task to obtain and interpret DS CLI exit codes.

About this task

Whenever you complete a transaction using the DS CLI single-shot mode or the script mode, an exit code is generated. However, no exit codes are generated when you use the DS CLI interactive mode, because you never leave the DS CLI session.

When you use the single-shot mode, an exit code is generated after each DS CLI command is fully processed. When you use the script mode, exit codes are only generated when the script exits the session. In script mode, you must interpret output for the status.

DS CLI exit codes provide more general reasons (than the error messages provide) why a CLI command transaction has failed. The following table lists the exit codes and their meanings.

Note: The operating system might generate other exit codes not listed in Table 9 if the DS CLI is stopped by another process. For example, code 143 for the Windows operating system.

Table 9. DS CLI exit codes

Code	Category	Description
0	Success	Specifies that the command is successfully processed.
2	Syntax error	Specifies that there is an error in the way that the command is presented (misaligned or wrong parameters) for processing.
3	Connection error	Specifies that there is a connectivity or protocol error.
4	Server error	Specifies that an error occurs during a function call to the application server.
5	Authentication error	Specifies that an error occurs during the authentication process.
6	Application error	Specifies that an error occurs due to a MetaProvider client application specific process.
63	Configuration error	Specifies that the CLI.CFG file is not found or accessible.
64	Configuration error	Specifies that the javaInstall variable was not provided in CLI.cfg.
65	Configuration error	Specifies that the javaClasspath variable was not provided in CLI.cfg.
66	Configuration error	Specifies that the format of the configuration file is not correct.

Complete the following steps to view, interpret, and use the DS CLI exit codes.

Procedure

1. (Script mode) Retrieve the most recent exit code. For a Windows operating system, use `%ERRORLEVEL%` to retrieve the most recent exit code. For a UNIX or Linux operating system, use `$?` to retrieve the most recent exit code.

   The following examples demonstrate the retrieval commands. The first part of the example shows the command that failed and the second part of the example shows the code to obtain the DS CLI exit code.

   Windows operating system

   ```
   C:\Program Files\ess\cli>dscli test
   CMMCI9013E Command: test was not found.
   Tip: Enter "help" for a list of available commands.

   C:\Program Files\ess\cli>echo %ERRORLEVEL%
   2
   ```

   UNIX or Linux operating system

   ```
   aix23 ->dscli test
    CMMCI9013E Command: test was not found.
   Tip: Enter "help" for a list of available commands.

   echo $?
   2
   ```

2. Use the previous table to interpret the value that is associated with the code and correct the command according to the exit code description.

Results

Processing that determines your next course of action

Based on the interpretation of the exit code value and the following processing description that is associated with a failed DS CLI transaction, you can determine your next course of action.

`Single-shot mode`

The following processing is associated with a single-shot mode transaction:

- All operations of the DS CLI transaction that can be processed are processed even though an error has occurred with one or more of the processed parameters that are associated with the transaction.
- A report on all successful completions is generated.
- A report on all failures is generated.

`Script mode`

The following processing is associated with a script mode transaction:

1. A DS CLI failure exit code is issued.
2. The script mode is automatically exited with no additional processing.

DS CLI operational limitations

Certain operational limitations are associated with the use of the DS CLI.

These limitations are described as follows:

- Volumes in the same volume space, logical subsystem (LSS), logical control unit (LCU), or address group cannot be of mixed type. They are either fixed block or count key data (CKD). The volume space is called an extent pool. An extent pool contains one or more ranks of a common storage type (fixed block or CKD).
- Logical subsystems cannot be created using the DS CLI. A fixed block LSS is automatically created when your first fixed block volume is assigned to the LSS address space. A fixed block LSS is automatically deleted when the last fixed block volume is removed from an LSS address space.

  Note: You can use DS CLI commands to create, modify, and delete LCUs, which are the CKD volume equivalent of a fixed block LSS.
- You must not start more than 100 DS CLI sessions simultaneously. Starting more than 100 DS CLI sessions simultaneously can result in connection problems.
- The maximum number of simultaneous DS CLI sessions is dependent on several execution environment factors such as the host hardware configuration, the vendor and version of the host OS, the vendor and version of the Java used by the DS CLI, the type and mix of DS CLI commands, LAN congestion, and the capacity, configuration, and current workload of your system. It is unlikely that you will experience problems using multiple DS CLI sessions, but under high stress conditions, you might be required to reduce the number of simultaneous DS CLI sessions on each host and the total number of DS CLI sessions to your system.
- The DS CLI is available from the HMC console with the following restrictions:
  - The DS CLI is only available in interactive mode from the fluxbox menu.
  - You cannot upgrade the DS CLI on the HMC except though newer releases of HMC software.
  - You cannot edit the existing DS CLI profile, or specify a different profile.
  - You cannot use commands that upload or download files, for example:
    - **`offloadauditlog`**
    - **`offloadfile`**

- **applykey** cannot be used with the -file parameter
- **setauthpol** cannot be used with the -action value of settruststore

Messages in the CLI and management console server

When you use the command-line interface and the management console, the applications provide messages regarding the application processes, status, and errors.

The user interfaces and the supporting software issue three types of messages. When you need to see the details about a DS CLI message, use the **helpmsg** command.

Informational messages
These messages are identified by the letter "I" at the end of the message identifier. They provide information about system activities as they take place. For example, an informational message might report that a volume was successfully created. No user action is necessary.

Warning messages
These messages are identified by the letter "W" at the end of the message identifier. They warn that user activated activities might have consequences that you do not anticipate. Warning messages normally provide the opportunity to continue an activity or to cancel it.

Error messages
These messages are identified by the letter "E" at the end of the message identifier. They indicate that an error has occurred. Refer to the explanations and recommended actions in this document to resolve the problem.

Chapter 3. CLI commands

About CLI commands

A description of the components and structure of a command-line interface command.

A command-line interface command consists of the following types of components, arranged in the following order:

1. The **command name**.
2. The **command flags** and **flag parameters.**
3. One or more **command parameters**, each followed by any **sub parameters** it might require.

The **command name** specifies the task that the command-line interface is to complete. For example, **lsarraysite** tells the command-line interface to list array sites, and **mklcu** tells the command-line interface to create a logical control unit.

Flags modify the command. They provide additional information that directs the command-line interface to complete the command task in a specific way. For example, the **-h** flag tells the command-line interface to display the reference page for the command. Some flags can be used with every command-line interface command. Others are specific to a command and are invalid when used with other commands. Flags are preceded by a hyphen (-), and can be followed immediately by a space and a flag parameter.

Flag parameters provide information that is required to implement the command modification that is specified by a flag. For example, the **-user** flag requires a *user_name* parameter, and the **-passwd** flag requires a *password* parameter. Flag parameters are variables, their value changes to meet your needs. Every user has a different user name and password. Not all flags require parameters. In this case, the flag itself provides all the information that is necessary. Some flag parameters are optional and might allow the use of multiple values. These values must be separated with a comma and no white space between the values. If you do not provide a parameter, then a default value is assumed. For example, you can specify *-v on*, or *-v off* to turn verbose mode on or off; but specifying *-v* only, results in an error.

The **command parameter** provides basic information that is necessary to complete the command task. When a command parameter is required, it is always the last component of the command; and it is not preceded by a flag. Some commands permit multiple command parameters with each parameter that is separated by a white space and not a comma (unlike flag parameters that allow multiple values). Some commands, like **lsuser**, do not require a command parameter because a default value of *all* is always assumed. For some commands, like **lsarraysite**, the command parameter is optional. If no value is provided, then a default value of *all* is assumed. If a value is provided, then the command-line interface lists information only about the array site or sites that are provided in the command parameter string.

In the following example, **lsrank** is the command name. **-dev** and **-l** are command parameters. *IBM.2107–75FA120* is the sub parameter for the **-dev** parameter, and *R1*, *R2*, and *R3* are a list of command parameters. Note that the banner is not listed for all examples provided in this document.

```
dscli> lsrank -dev IBM.2107-75FA120 -l R1 R2 R3
```

Understanding the syntax diagrams

A syntax diagram uses symbols to represent the elements of a command and to specify the rules for using these elements.

© Copyright IBM Corp. 2004, 2019

Syntax diagrams

Main path line

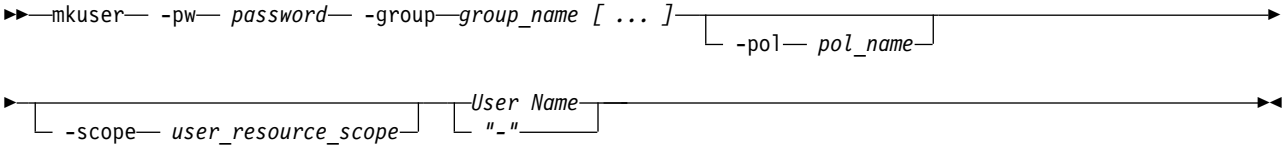
Begins on the left with double arrowheads (>>) and ends on the right with two arrowheads facing each other (><). If a diagram is longer than one line, each line to be continued ends with a single arrowhead (>) and the next line begins with a single arrowhead. Read the diagrams from left-to-right, top-to-bottom, following the main path line.

Keyword



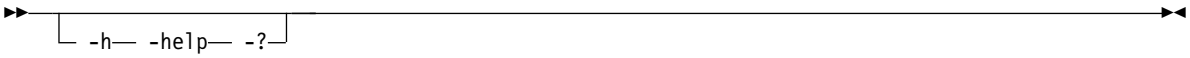
Represents the name of a command, parameter, or argument. A keyword is not in italics. Spell a keyword exactly as it is shown in the syntax diagram.

Required keywords



Indicates the parameters or arguments you must specify for the command. Required keywords appear on the main path line. Mutually exclusive required keywords are stacked vertically.

Optional keywords



Indicates the parameters or arguments you can choose to specify for the command. Optional keywords appear below the main path line. Mutually exclusive optional keywords are stacked vertically.

Variable

Represents the value you need to supply for a parameter or argument, such as a file name, user name, or password. Variables are in italics.

Special characters

- (minus) or / (slash) sign
Parameters are prefixed with a - (minus) sign. Parameters define the action of a command or modify the operation of a command. You can use multiple parameters, followed by variables, when you issue a command.

[ ] square brackets
Optional values are enclosed in square brackets.

{ } braces
Required or expected values are enclosed in braces.

| vertical bar
A vertical bar indicates that you have a choice between two or more options or arguments.

For example, [a | b] indicates that you can choose a, b, or nothing. Similarly, { a | b } indicates that you must choose either a or b.

... ellipsis
An ellipsis signifies the values that can be repeated on the command line or multiple values or arguments.

- dash A dash indicates that, as an alternative to entering the parameter, a value or values are supplied from stdin. stdin varies depending on your settings and is available when you are using single-shot or script mode. This option is not available when using interactive mode.

Common command flags

You can use these flags with any command-line interface command.

Parameters	Description
`-p` on | off	Turns paging on or off. Displays 24 rows at a time unless used with the `-r` flag. The default is off in single-shot mode and on in interactive mode. You can page up or down by pressing any key. You can specify *-p on* or *-p off* to turn the paging on or off; but if you specify only *-p*, you get an error. **Note:** You can only use this flag with the ls type (for example, lsuser, lskey, lsserver) commands and the help (setoutput) command.
`-r` *number*	Specifies the number of rows (1 - 100) per page. This flag is valid only when the `-p` flag is set to on. The default value is 24 rows. **Note:** You can only use this flag with the ls type (for example, lsuser, lskey, lsserver) commands and the help (setoutput) command.
`-fmt` xml | stanza | delim | default	**xml** Sets the output format to XML. **Note:** You can use this option only with list (for example, lsuser, lskey, lsserver) commands **stanza** Sets the output format to stanza. **Note:** You can use this option only with list (for example, lsuser, lskey, lsserver) commands **delim** Sets the output format to a table. You must set the column delimiter to a single character with the -delim flag. **Note:** You can use this option only with list (for example, lsuser, lskey, lsserver) commands **default** Sets the output to a space-separated plain text table. **Note:** You can use this option only with list (for example, lsuser, lskey, lsserver) commands
`-delim` *char*	Sets the output to delimited output and the delimiter to the single character *char*. You must enclose *char* in single or double quotation marks if the character is a shell metacharacter (such as * or `\t`). If *char* is not specified, the CLI program returns a syntax error. A blank space, even when it is enclosed within quotation marks, is not a valid character as a delimiter. **Note:** You can use this option only with list (for example, lsuser, lskey, lsserver) commands
`-hdr` on | off	Turns the header on or off. The default is on. You can specify *-hdr on* or *-hdr off* to turn the header on or off; but if you specify only *-hdr*, you get an error.

Parameters	Description
`-bnr on | off`	Turns the banner on or off. The default is on. You can specify *-bnr on* or *-bnr off* to turn the banner on or off; but if you specify only *-bnr*, you get an error.
`-v on | off`	Turns verbose mode on or off. The default is off. You can specify *-v on* or *-v off* to turn verbose mode on or off; but if you specify only *-v*, you get an error.
`-fullid`	Provides fully qualified IDs, which include the storage image ID, for every ID that is displayed in the command output. **Note:** You can use this command flag only with the list (for example, lsioport, lskey) and show (for example, showsu, showlss) commands.
`-help | -h | -?`	Displays a detailed description of the specified command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.
`-s`	Displays the short output for the command, which is usually a single column of the resource IDs of the objects specified by the command. For detailed information, see the **-s** parameter for that command. **Notes:** 1. You can only use this flag with the help (setoutput) command and the ls type (for example, lsuser, lsserver) commands, except for lskey. 2. You cannot use the `-s` and `-l` parameters together.
`-l`	Displays the long output for the command, which is the default output plus additional columns of data to the right of the default columns. For detailed information for a specific command, see the **-l** parameter for that command. **Notes:** 1. You can only use this flag with the help (setoutput) command and the ls type (for example, lsuser, lsserver) commands, except for lskey. 2. You cannot use the `-s` and `-l` parameters together.

Application key commands

Application key commands activate and query the licenses on your storage system.

The following application key commands are available:

applykey
Activates the licenses on your storage system.

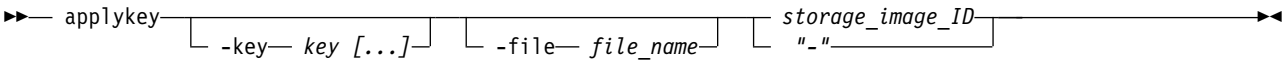
lskey Generates a report that displays the type of license keys that are installed and are available for use by the specified storage system.

applykey

The **applykey** command applies the licensed machine code (LMC) activation keys for a storage server.

You can enter the LMC keys manually, or you can import the keys from an XML file, which you must download.

Note: The **applykey** command cannot be run from the Embedded DS CLI window.



Parameters

-key *key [...]*
(Optional) Specifies the LMC key. The optional ellipsis (...) signifies that you can specify multiple keys. To specify multiple keys, enter a comma between each key. Do not include a blank space between each key.

This parameter is required if the **-file** parameter is not specified.

-file *file_name*
(Optional) Specifies the file name of the LMC activation key file.

This parameter is required if the **-key** parameter is not specified.

storage_image_ID | **-**
(Required) Specifies the storage image ID where the LMC activation key file is imported. The ID includes manufacturer, machine type, and serial number.

If you use the dash (-), the specified value is read from standard input.

Example

Invoking the applykey command (importing the key)

```
dscli> applykey -file keys.xml IBM.2107-75FA120
```

Example

Invoking the applykey command (manually entering the key)

```
dscli> applykey -key DA67-7D5C-4B98-2E1B-C7FA-8F85-BD8A-31AD
IBM.2107-75FA120
```

lskey

The **lskey** command displays the type of LMC activation keys that are installed and are available for use by the storage system.

The **lskey** command displays only the keys that are installed. Refer to the *IBM DS8000 Introduction and Planning Guide* for more information on how to choose which keys are needed and how to acquire them.



Parameters

storage_image_ID | **-**
(Required) Specifies the storage image ID for which to view a list of activated features. The ID includes manufacturer, type, and serial number.

If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format for clarity. The actual reports do not display as tables.

The following table shows example activation keys. Some activation keys are not listed in the example. The **lskey** command displays only the keys that are installed.

Invoking the lskey command to display a list of the installed LMC activation keys.

```
dscli> lskey IBM.2107-75FA120
```

The resulting output

Activation key	Authorization level (TB)	Scope
Parallel access volumes (PAV)	On	CKD
Point in time copy (PTC)	On	All
Global Mirror (GM)	25	All
HyperPAV	On	CKD
Metro/Global Mirror (MGM)	25	All
Metro Mirror (MM)	25	All
Remote mirror and copy (RMC)	25	All
Remote mirror for z/OS (RMZ)	25.1	CKD
Operating Environment (OEL)	45	All
IBM Easy Tier	On	All
IBM FlashCopy SE	25	All
High Performance FICON for System z (zHPF)	On	CKD
RMZ Resync	25	CKD

Report field definitions

Activation key
Indicates the type of LMC activation key that is activated for the storage image.

Authorization Level (TB)
Indicates the capacity of the specified license feature. The quantity is displayed in terabytes (TB). One of the following values is displayed:

- Value in terabytes
- **On** if the license is for the maximum capacity, or **Off** if the license is for zero capacity

Scope
Indicates the storage type for the designated license: fixed block (FB), count key data (CKD), or All. Parallel access volumes, Remote Mirror for z/OS, and HyperPAV display only the values CKD or All.

Cloud commands

Cloud server commands are available to help manage DS8000 resources.

lscloudserver
Displays a list of cloud server resources on the storage system.

managecloudserver
Adds one or more TS7700 cloud server resources to the node.

mkcloudserver
Creates a configuration entry for a cloud server provider.

rmcloudserver
Removes a cloud server that is running on the storage system.

lscloudserver

The **lscloudserver** command displays a list of clouds on the storage system.



Parameters

-type swift | swift-keystone | TS7700 | aws-s3 | ibmcos | s3
(Required) Displays the type of cloud. The following options are supported:

swift
Specifies that the DS8000 system uses OpenStack Swift cloud computing system to encrypt authentication credentials and connect to a cloud storage target.

swift-keystone
Specifies that the DS8000 system uses Swift-Keystone to encrypt authentication credentials and connect to a cloud storage target. Authentication is done by using root or system certificates with either Secure Sockets Layer or Transport Layer Security, SSL/TLS.

TS7700
Specifies that the DS8000 system authenticates and stores data directly on the TS7700 system. Authentication credentials such as a user name, password, or certificate are not required. Instead, the DS8000 system uses primary and alternate IP addresses to authenticate and verify connection to the TS7700 system.

aws-s3
Specifies that Amazon Simple Storage Service (Amazon S3) is to allow the DS8000 system to authenticate and connect to S3 storage.

ibmcos
Specifies IBM Cloud Object Storage (COS) for data protection through backup and recovery.

s3 Specifies that s3 allows the DS8000 system to authenticate and store data on any other s3 compatible cloud storage target.

Example

```
dscli> lscloudserver -l
```

The resulting output

name	node	type	tenant	location	user	endpoint	IP address	nossl
ibmcloud	0	swift	mytest04	-	mrproduct	https://hostname:port/	-	false
ibmcloud	1	swift	mytest04	-	mrproduct	https://hostname:port/	-	false

Report field definitions

name*
Displays the cloud name.

node*
Displays the node that is associated with the cloud.

type*
Displays the cloud type.

tenant+
Displays the tenant name of the cloud.

location+
Displays the location for the cloud.

user+
Displays the user name of the cloud.

endpoint*
Displays the address of the cloud.

IP address*
Displays the IPv4 address of the cloud for type TS7700 only.

nossl+
Displays the setting (true or false) of this parameter. If the value is true, the SSL protocol is not used when transparent cloud tiering communicates with the specified cloud. Otherwise, the SSL protocol is used.

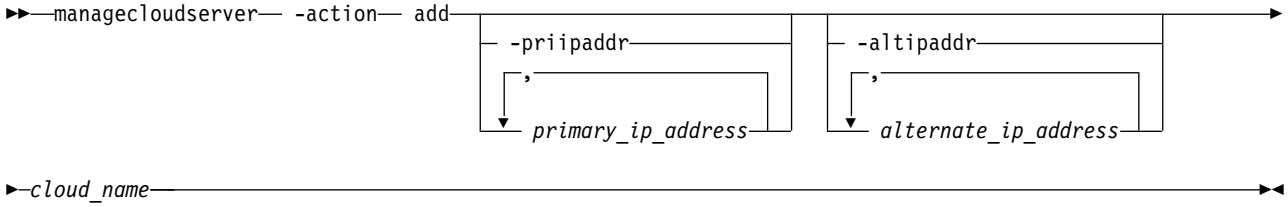
Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

managecloudserver

The **managecloudserver** command adds one or more TS7700 clouds to the node.



Parameters

-action add
(Required) Specifies the action type. Only the **add** parameter is supported.

-priipaddr *primary_ip_address*
(Optional) Specifies up to two primary comma-separated IPv4 addresses.

-altipaddr *alternate_ip_address*
(Optional) Specifies up to two alternate comma-separated IPv4 addresses.

cloud_name
(Required) Specifies the name of the cloud. Only one cloud is supported.

Example

Invoking the managecloudserver command

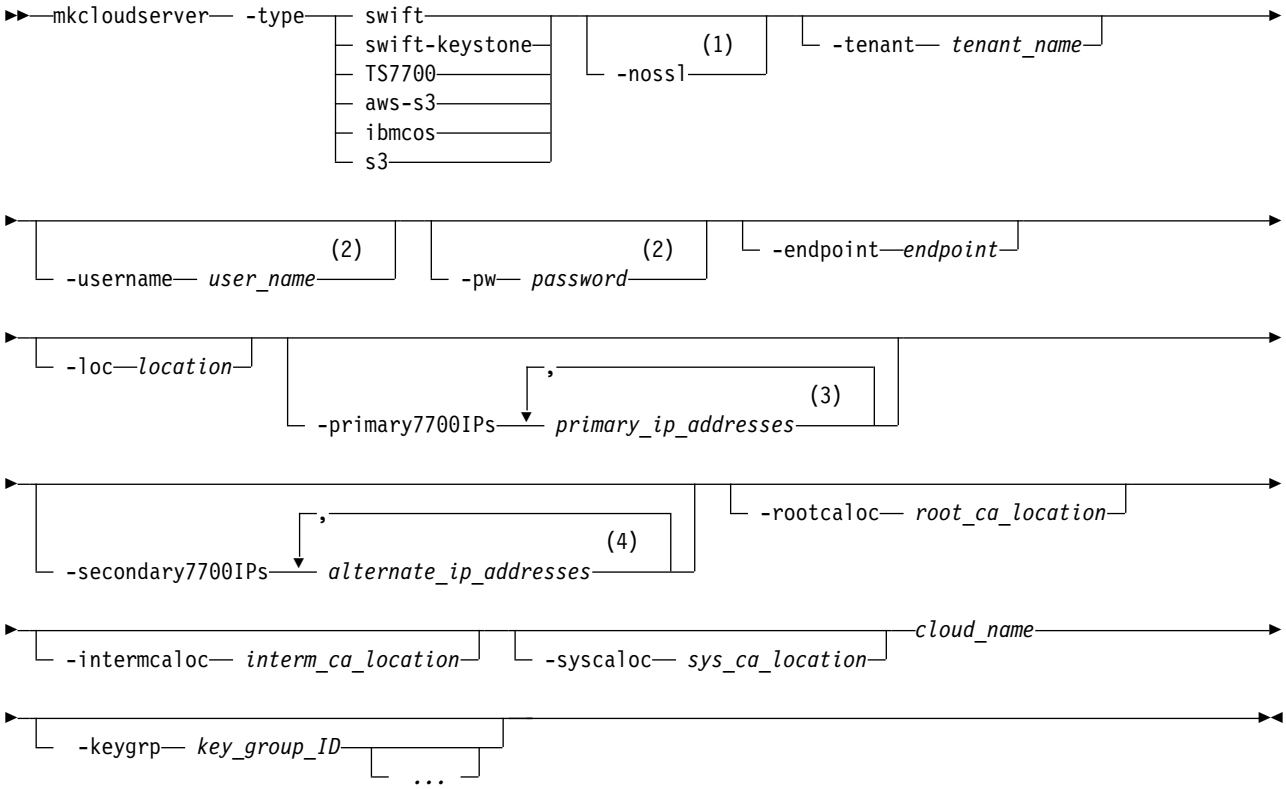
```
dscli> managecloudserver -action add -priipaddr 20.18.153.26,20.18.153.27
 -altipaddr 20.18.152.26,20.18.152.27 ibmcloud
```

The resulting output

```
Date/Time: February 20, 2018 2:58:50 AM MST IBM DSCLI Version: 0.0.0.0 DS: -
CMUC00501I managecloudserver: The entered cloud ibmcloud was operated
 successfully on node 0.
CMUC00501I managecloudserver: The entered cloud ibmcloud was operated
 successfully on node 1.
```

mkcloudserver

The **mkcloudserver** command creates a configuration entry for a cloud.



Notes:

1 The **nossl** parameter is required when the cloud type is TS7700 or swift.

2 The parameters **tenant, username, pw, endpoint, rootcaloc, intermcaloc,** and **syscaloc** cannot be used when the cloud type is TS7700.

3 The **primary7700IPs** parameter is required when the cloud type is TS7700.

4 The **secondary7700IPs** parameter is required when the cloud type is TS7700.

Parameters

-type swift | swift-keystone | TS7700 | aws-s3 | ibmcos | s3
(Required) Displays the type of cloud. The following options are supported:

swift
Specifies that the DS8000 system uses OpenStack Swift cloud computing system to encrypt authentication credentials and connect to a cloud storage target.

swift-keystone
Specifies that the DS8000 system uses Swift-Keystone to encrypt authentication credentials and connect to a cloud storage target. Authentication is done using root or system certificates with either Secure Sockets Layer or Transport Layer Security, SSL/TLS.

TS7700
Specifies that the DS8000 system authenticates and stores data directly on the TS7700 system. Authentication credentials such as a user name, password, or certificate are not required. Instead, the DS8000 system uses primary and alternate IP addresses to authenticate and verify connection to the TS7700 system.

aws-s3
Specifies that Amazon Simple Storage Service (Amazon S3) is to allow the DS8000 system to authenticate and connect to S3 storage.

ibmcos
Specifies IBM Cloud™ Object Storage (COS) for data protection through backup and recovery.

s3 Specifies that s3 allows the DS8000 system to authenticate and store data on any other s3 compatible cloud storage target.

-nossl
(Optional) Specifies that communication between the DS8000 and the specified cloud does not require the SSL protocol.

Note: The **-nossl** parameter is required when you specify the **-type** parameter with either option `TS7700` or `swift`.

-tenant *tenant_name*
(Optional) Specifies the tenant name that you use to connect with the cloud.

-username *user_name*
(Optional) Specifies the user name that you use to connect with the cloud.

Note: The **-username** parameter cannot be used with type TS7700.

The **-username** parameter is required when you specify the **-type** parameter with the `swift`, `swift-keystone`, `aws-s3`, `s3`, or `ibmcos` option.

-pw *password*
(Optional) Specifies the password that you use to connect with the cloud.

Note: The **-pw** parameter cannot be used with type TS7700.

The **-pw** parameter is required when you specify the **-type** parameter with the `swift`, `swift-keystone`, `aws-s3`, `s3`, or `ibmcos` option.

-endpoint *endpoint*
(Optional) Specifies the location, or URL, of the cloud.

Note: The **-endpoint** parameter cannot be used with type TS7700.

The **-endpoint** parameter is required when you specify the **-type** parameter with the `swift`, `swift-keystone`, `aws-s3`, or `ibmcos` option.

-loc *location*
(Optional) Specifies the detailed location or vault template of the cloud depending on the -type parameter.

Note: The **-loc** parameter is valid with type `ibmcos` or `aws-s3`, only.

ibmcos
The location parameter specifies the name of the user generated vault template created on the IBM Cloud Object Storage Manager.

aws-s3
The location parameter specifies the AWS Region as defined by the endpoint of the Amazon S3 service.

-primary7700IPs *primary_ip_address*
(Required for TS7700 only) Specifies the primary IP address of the cloud.

Note: You can specify up to four primary IP addresses (IPv4 is supported). Separate the primary IP addresses with a comma.

-secondary7700IPs *alternate_ip_address*
(Optional for TS7700 only) Specifies the alternate IP address of the cloud.

Note: You can specify up to four alternate IP addresses (IPv4 is supported). Separate the alternate IP addresses with a comma.

-rootcaloc *root_ca_location*
(Optional) Specifies the path of the root certification file on your computer. The certificate determines which protocol (such as Transport Layer Security, TLS) that is used to communicate with the cloud.

Multiple clouds can share the root certificate if a root certificate exists that is defined to the system. (Root and intermediate certificates are currently the same.) If a root certificate does not exist in the storage system, the **-rootcaloc** parameter is required.

Note: The **-rootcaloc** parameter cannot be used with type TS7700. The **-rootcaloc** parameter is required if a root certificate does not exist for the storage system.

-intermcaloc *interm_ca_location*
(Optional) Specifies the path of the intermediate certification file on your computer.

Note: The **-intermcaloc** parameter cannot be used with type TS7700.

-syscaloc *sys_ca_location*
(Optional) Specifies the path of the system certification file on your computer.

The certificate determines which protocol (such as TLS) to use to communicate with the cloud. Multiple clouds can share root certificate. Each cloud has a respective system certificate.

Note: The **-syscaloc** parameter cannot be used with type TS7700.

cloud_name
(Required) Specifies the name of the cloud. DS8000 supports one cloud.

-keygrp *key_group_ID*
(Optional) Specifies a key group that is configured for encryption of transparent cloud tiering, which ensures that data is encrypted while it is transferred to cloud storage.

Example

```
dscli> mkcloudserver -type swift -tenant mytest04 -username mrproduct -pw ProdUct02
-endpoint "https://hostname:port/"  -rootcaloc "/tmp/ssl_cacert.pem"
-syscaloc "/tmp/ssl_cert.pem" ibmcloud
```

Output

```
CMUC00505I mkcloudserver: The entered cloud ibmcloud was created
 successfully on node 0.

CMUC00505I mkcloudserver: The entered cloud ibmcloud was created
 successfully on node 1.
```

rmcloudserver

The **rmcloudserver** command removes a cloud that is running on the storage system.



Parameters

-quiet

(Optional) Turns off the confirmation prompt for this command.

cloud_name

(Required) Specifies the name of the cloud to be deleted. DS8000 supports one cloud.

Example

Invoking the rmcloudserver command

```
dscli> rmcloudserver ibmCloud
```

The resulting output

```
CMUC00503W Are you sure you want to delete the cloud
      ibmcloud? [y/n]
CMUC00504I The entered cloud ibmcloud was deleted
      successfully from node 1.
```

Cluster commands

Cluster commands are available to help manage a group of host systems called a cluster.

chcluster
Modifies the configuration of a cluster.

lscluster
Displays a list of clusters and status information for each cluster in the list.

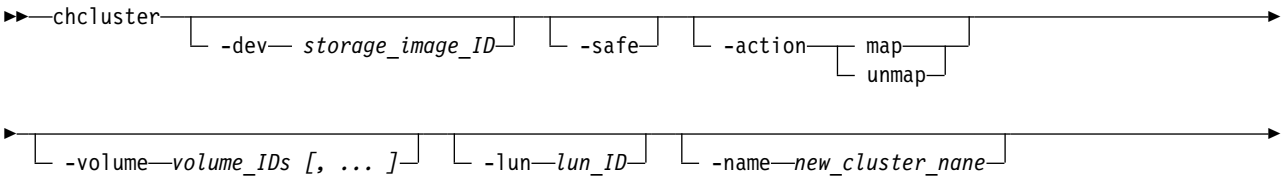
mkcluster
Configures a cluster for the storage system.

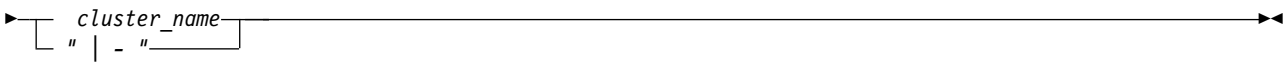
rmcluster
Removes a cluster from a storage system.

showcluster
Displays detailed properties of an individual cluster.

chcluster

The **chcluster** command modifies the configuration of a cluster.





Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-safe
(Optional) Enables checks for the safe removal of one or more specified volumes from the host.

-action
(Optional) Specifies the type of action for the specified volumes.

map
Specifies that the volumes are added to the cluster.

unmap
Specifies that the volumes are removed from the cluster.

Note: The **chcluster** command fails if you specify the **-volume** parameter without the **-action** parameter.

-volume *volume_ID [...]*
(Optional) Specifies an array of one or more volume IDs or volume ID ranges to be added or removed from the specified host. A volume ID range is defined by two volume IDs that are separated by a hyphen. Multiple volume IDs or volume ID ranges must be separated by a comma. The volume IDs are entered as hexadecimal numbers in the range 0x0000-0xFEFF. Example: 0100-010F, 0180-018F,0120.

Notes:

1. For SCSI MAP 256, the array or ranges cannot exceed 256 volume ID entries. Otherwise, up to 64,384 entries are allowed.
2. The **chcluster** command fails if you specify the **-volume** parameter without the **-action** parameter.

-lun *lun_ID*
(Optional) Specifies the LUN ID in hexadecimal value (00 - FF) that is mapped to the specified volume ID for a SCSI-MAP256 type cluster. If you specify multiple volume IDs with the **-volume** parameter, LUN IDs are assigned consecutively in incremental order.

Note: This parameter is only valid for a SCSI-MAP 256 type cluster. If this parameter is specified for any other type of cluster, the command fails.

-name *new_cluster_name*
(Optional) Specifies the new name you assign to the cluster system. The name is limited to 32-byte or 16-double byte characters. If you assign a new cluster name with a blank space, enclose the name with double quotation marks.

cluster_name | **-**
(Required) Specifies the cluster name to be modified. If the cluster name you specify includes a blank space, enclose the name with double quotation marks.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the chcluster

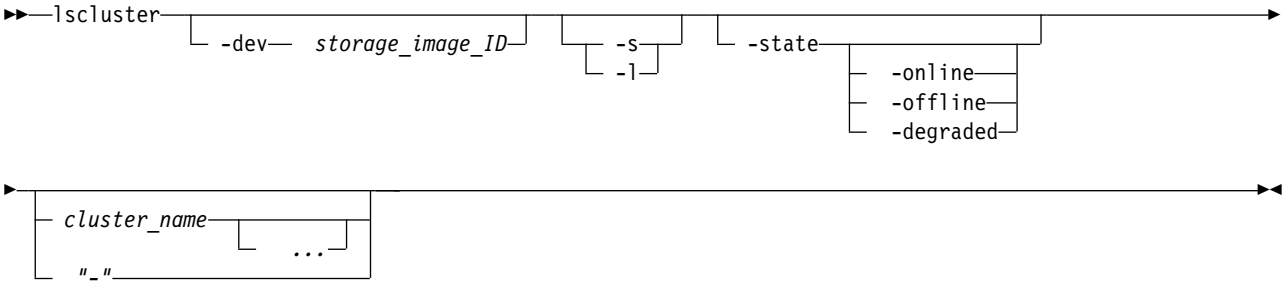
```
dscli> chcluster -dev IBM.2107-75FA120 -name cluster_1 cluster_2
```

The resulting output

```
Cluster cluster_2 successfully modified.
```

lscluster

The **lscluster** command displays a list of cluster and status information for each cluster in the list.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-s
(Optional) Displays the cluster name only. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output and user-defined description for each cluster in the list. You cannot use the **-l** and the **-s** parameters together.

-state
(Optional) Displays all clusters with the following specified state.

online
All hosts that the cluster contains are online.

offline
All hosts that the cluster contains are offline.

degraded
Not all hosts, but at least one, that the cluster contains are online.

cluster_name ... | -
(Optional) Displays the cluster information for one or more cluster names specified. The ellipsis (...) indicates that you can specify multiple values. Multiple names must be separated with a white space between each name or ID. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

```
dscli> lscluster -dev IBM.2107-75FA120 -l
```

The resulting output

Name	Type	State	AddrMode	Volumes	Hosts
Cluster_1	Win2003	Online	VMWare	8	5
Cluster_2	-	Offline	unknown	6	0
Cluster_3	Mix	Degraded	scsimap256	0	8

Report field definitions

Name*
Displays the name of the cluster. The name is limited to 32-byte or 16- double-byte characters.

Type
Displays the host type that is defined by the operating system and platform.

If the hosts assigned to the specified cluster have a different host type, the output displays "multiple". If are no hosts are assigned to the specified cluster, the output displays "-".

State
Displays the state of the cluster. One of the following values displays.

online
Indicates that all hosts that the cluster contains are online.

offlne
Indicates that all hosts that the cluster contains are offline.

degraded
Indicates that at least one host that the cluster contains is online.

AddrMode+
Indicates the method that is used by the cluster to discover the LUNs that are available to its host ports. One of the following values display.

scsimask
A SCSI mask volume group. This option is available if the host adapter supports 4-byte LUN addresses.

scsimap256
A SCSI-MAP 256 volume group.

os400mask
An OS/400 mask volume group. The IBM i host system typically uses fixed block volumes of 520-byte logical block size. This option is available only if the host adapter supports 4-byte LUN addresses.

unknown
If no host that is associated with the cluster is available, the output displays "unknown".

Volumes+
Indicates the number of public volumes that are mapped to the specified cluster.

Hosts+
Indicates the number of hosts that are mapped to the specified cluster.

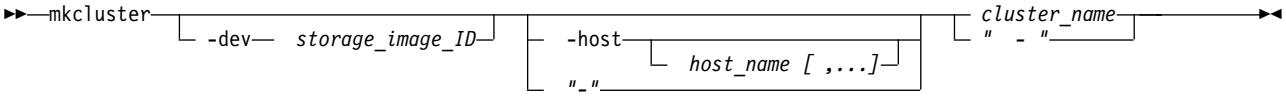
Key:

* Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

mkcluster

The **mkcluster** command configures a cluster for the storage system. A cluster is defined as a group of host systems associated with the storage system.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-host *host_name [,...]*
(Optional) Specifies one or more host names that are assigned to the cluster. Separate multiple host names by a comma. If the host name includes a blank space, enclose the name with double quotation marks.

cluster_name | **-**
(Required) Specifies the name that you assigned to the cluster. The name is limited to 32-byte or 16 double byte characters. If the cluster name you specify includes a blank space, enclose the name with double quotation marks.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the mkcluster command

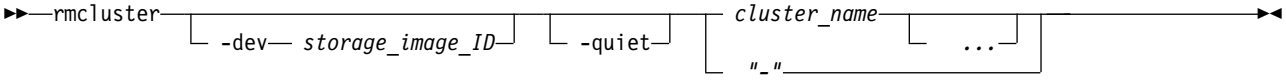
```
dscli> mkcluster cluster_1
```

The resulting output

```
Cluster cluster_1 successfully created.
```

rmcluster

The **rmcluster** command removes a cluster from a storage system.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-quiet
(Optional) Turns off the confirmation prompt for this command

cluster_name ... | -
(Required) Specifies the cluster name to be removed. The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmcluster

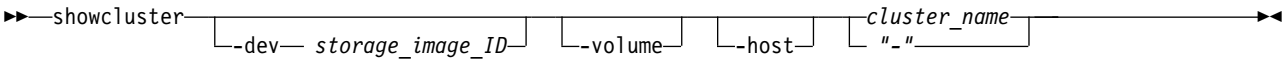
```
dscli> rmcluster -dev IBM.2107-75FA120 cluster_1
```

The resulting output

```
Are you sure you want to delete Cluster cluster_1? y/n Y
Cluster cluster_1 successfully deleted.
```

showcluster

The **showcluster** command displays detailed properties of an individual cluster.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-volume
(Optional) Displays the volume and LUN ID map for the specified cluster.

-host
(Optional) Displays the host name for the specified cluster.

cluster_name | -
(Required) Displays the name of the cluster. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 1

```
dscli> showcluster Cluster_1
```

The resulting output

```
Name                Cluster_1
Type                Win2003
State               online
AddrMode            scsimask
Volumes             3
Hosts               3
```

Report field definitions

Name
Displays the name of the cluster. The name is limited to 32-byte or 16-double-byte characters.

Type
Displays the host type that is defined by the operating system and platform.

If the hosts under the specified cluster have a different host type, the output displays "multiple". If there are no hosts that are associated with the specified cluster, the output displays "-".

State
Displays the state of the cluster. One of the following values displays.

online
Indicates that all hosts that the cluster contains are online.

offlne
Indicates that all hosts that the cluster contains are offline.

degraded
Indicates that at least one host that the cluster contains is online.

AddrMode
Displays the method that is used by the cluster to discover the LUNs that are available to its host ports. One of the following values display.

scsimask
A SCSI mask volume group. This option is available if the host adapter supports 4-byte LUN addresses.

scsimap256
A SCSI-MAP 256 volume group.

os400mask
An OS/400® mask volume group. The IBM i host system typically uses fixed block volumes of 520-byte logical block size. This option is available only if the host adapter supports 4-byte LUN addresses.

unknown
If no host is available under the cluster, the output displays "unknown".

Volumes
Displays the number of volumes that are mapped to this cluster.

Host
Displays the number of hosts that are mapped to this cluster.

Example 2

Invoking the showcluster command with the volume parameter specified

```
dscli> showcluster -volume Cluster_1
```

The resulting output

```
Name                 Cluster_1
Type                 Win2003
State                online
AddrMode             scsimask
Volumes              3
Hosts                3
================Volumes===========
ID              Lun
==================================
000D            8D
000E            8E
000F            8F
```

When you specify the **-volume** parameter, the Volume ID and Lun ID also display.

ID Displays the ID of the specified volume.

Lun
Displays the Lun ID that is mapped to the volume ID.

Example 3

Invoking the showcluster command with the host parameter specified

```
dscli> showcluster -host Cluster_1
```

The resulting output

```
Name               Cluster_1
Type               Win2003
State              Online
AddrMode           scsimask
Volumes            3
Hosts              3
==============Hosts================
Name
==================================
Host_1
Host_2
Host_3
```

When you specify the **-host** parameter, the name parameter also displays.

Name
Displays the name that is assigned to the host.

Copy Services commands

For reference, the various Copy Services commands are listed with descriptions of their functions.

You can use the following Copy Services commands to manage Copy Services tasks.

FlashCopy commands

Commands that are used to configure FlashCopy relationships and to display FlashCopy information are referenced.

The following FlashCopy commands are available:

commitflash
Completes a partially formed Global Mirror consistency group. It is used as part of the recovery from a disaster.

resyncflash
Creates a point-in-time copy of an existing FlashCopy pair that was established with the **-record** and **-persist** parameters. The **resyncflash** command only copies the parts of the volume that have changed since the last point in time copy.

lsflash
Generates a report that displays a list of FlashCopy relationships and the status information for each FlashCopy relationship in the list.

mkflash
Initiates a point-in-time copy from source volumes to target volumes.

reverseflash
Reverses the FlashCopy relationship.

revertflash
Restores the former Global Mirror consistency group from one that is currently forming. It is used as part of the recovery from a disaster.

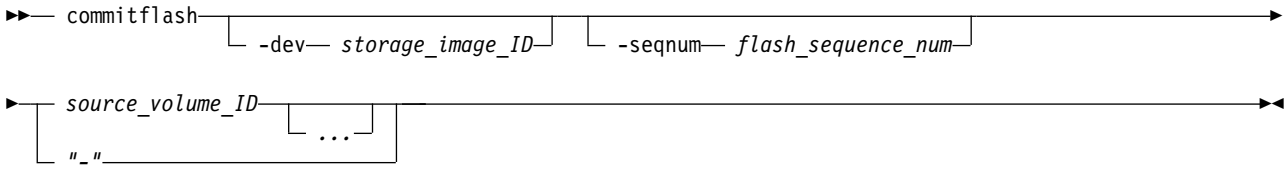
rmflash
Removes a relationship between FlashCopy volume pairs.

unfreezeflash
Resets a FlashCopy consistency group that was previously established with the **-freeze** parameter when the **mkflash** or **resyncflash** commands were issued.

setflashrevertible
Modifies a FlashCopy volume pair that is part of a Global Mirror relationship to revertible. The revertible feature allows data to be committed to the target to form a new consistency group, or restored back to the last consistency group.

commitflash

The **commitflash** command is used as part of the recovery from a disaster scenario to complete a partially formed Global Mirror consistency group.



Parameters

The following transactions must be completed before you can issue the **commitflash** command:

1. Issue the **mkflash** command with the **-record** and **-persist** parameters specified to establish the FlashCopy volume pair relationship.
2. Issue the **setflashrevertible** command on the FlashCopy volume pair.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for all source volumes, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-seqnum *flash_sequence_number*
(Optional) When a FlashCopy sequence number is specified, the commit operation is performed only on those relationships that are associated with the specified number.

This parameter is not supported for machine type 2105.

Example: 0010

source_volume_ID ... | **-**
(Required) Specifies the source volumes for which FlashCopy relationships are to be committed. The chosen FlashCopy pair is the one established or modified with the **-record** parameter. This parameter accepts fully qualified volume IDs, which includes storage image IDs, or a shortened version without storage image IDs if either the **-dev** parameter is specified or you specify a value for the *devid* variable in your profile file. You must separate multiple source volume IDs with spaces.

The volume ID is a 32-bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ* where:

X Specifies the address group, 0-F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified volume ID: IBM.2107-75FA120/0001

Example of a shortened version: 0001

Example of multiple IDs: 0001 0003 0008

Example

Invoking the commitflash command

```
dscli> commitflash -dev IBM.2107-75FA120 0100
```

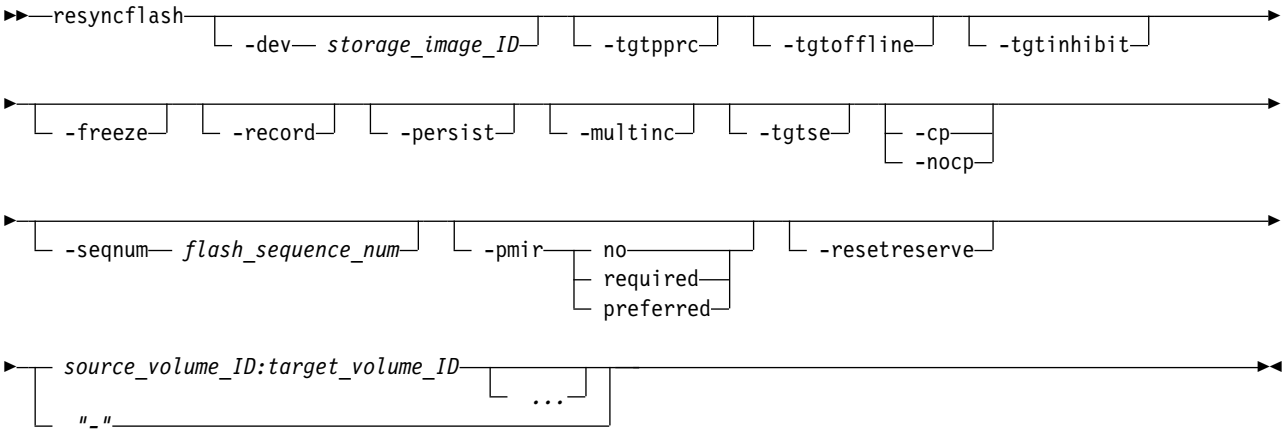
The resulting output

```
FlashCopy volume pair for source 0100 successfully committed.
```

resyncflash

The **resyncflash** command is a point-in-time copy of an existing FlashCopy pair that was established with the **-record** and **-persist** parameters.

The **resyncflash** command copies only the parts of the volume that changed since the last point-in-time copy.



Parameters

When a FlashCopy pair is established with the **-record** and **-persist** parameters, the pair initially synchronizes. A record of all host write operations to the source is maintained in the source volumes. When the **resyncflash** command is specified on the FlashCopy pair, the new data that is written to the source is copied to the target. The parameters that you specify in this command replace the parameters that you previously specified for the existing relationship. To keep the **-record** and **-persist** parameters, you must specify these parameters in the **resyncflash** command.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the source and target volume or do not set the *devid* variable in your profile or through the **setenv** command. The

storage image ID is also required if HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-tgtpprc
(Optional) Allows the FlashCopy target volume to be a remote mirror and copy source volume.

-tgtoffline
(Optional) Rejects the **resyncflash** command if the target volume is online for host system access. This parameter applies only to count key data (CKD) volumes.

-tgtinhibit
(Optional) Prevents host system write operations to the target volume while the FlashCopy relationship exists.

-freeze
(Optional) Starts the *queue full* condition for the source volume. During the *queue full* condition, the source volume reports a status of *long busy*. All writes to the source volume are queued by the host and are written after the *queue full* condition is reset. The *queue full* condition is reset by an *extended long busy timeout* condition. The timeout condition affects all FlashCopy source volumes that are contained within a respective logical subsystem and that are established or modified with this parameter.

Note: Use the **chlss** and the **chlcu** commands to modify the extended long busy timeout setting.

-record
(Optional) This parameter, without the **-multinc** parameter, creates a type 1 incremental FlashCopy relationship. The type 1 FlashCopy records data changes on both the source and target volumes of the FlashCopy pair.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

Select this parameter when you create an initial FlashCopy relationship that you later want to use with the **resyncflash** or **reverseflash** command. If the **-multinc** parameter is not selected, you can also use the **setflashrevertible** command.

When you select the -**record** parameter, the –**persist** parameter is automatically selected.

-persist
(Optional) Creates a persistent FlashCopy relationship in which the relationship remains after the copy completes and remains indefinitely until a **rmflash** command is issued against the FlashCopy pair. If this parameter is not specified, a normal FlashCopy relationship is created and is automatically removed after the copy completes.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

See the **-record** parameter for a description of a type 1 incremental FlashCopy relationship and the **-multinc** parameter for a description of a type 2 incremental relationship.

When you select either the **-record** or the **-multinc** parameter, the **persist** parameter is automatically selected.

-multinc
(Optional) Creates a type 2 incremental FlashCopy relationship. The type 2 FlashCopy records data changes only on the target volume of the FlashCopy pair.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

The type 2 FlashCopy allows for more than one incremental FlashCopy relationship from the same source volume. However, because the change recording is maintained only on the target volume, the type 2 FlashCopy can cause a performance impact as more type 2 FlashCopy relationships are added.

Select this parameter when you create multiple FlashCopy volume pairs with the same source volume that you want to use with the **resyncflash** and **reverseflash** commands. However, FlashCopy pairs that are established with this modified recording method cannot be used with the **setflashrevertible** command.

When you select the **-multinc** parameter, the **-persist** and **-record** parameters are automatically selected.

Note: Incrementing multiple relationships with the same source volume on a single command line (as shown below) is not currently supported:

```
resyncflash -multinc 0000-00ff:0200-02ff 0000-00ff:0400-04ff 0000-00ff:0600-06ff
```

Instead, you must issue the **resyncflash** commands separately. For example:

```
resyncflash -multinc 0000-00ff:0200-02ff
resyncflash -multinc 0000-00ff:0400-04ff
resyncflash -multinc 0000-00ff:0600-06ff
```

-tgtse
(Optional) Specifies that the target volume that you are designating for a FlashCopy relationship might be a space-efficient logical volume. An error message is generated if the target volume that you are using to create the FlashCopy relationship is a space-efficient volume and you do not specify this parameter.

-nocp
(Optional) Inhibits a background copy. Data is copied from the source volume to the target volume only if a track on the source volume is modified. The FlashCopy volume pair relationship remains indefinitely until it is broken by a **rmflash** command or until all tracks on the source volume are modified.

When the **-tgtse** parameter is specified and the **-cp** or the **-nocp** parameters are specified, the no background copy behavior is the default.

Note: You cannot use the **-nocp** parameter with the **-cp** parameter in the same command.

-cp
(Optional) Starts a background copy. When the **-tgtse** parameter is not specified and the **-cp** or the **-nocp** parameters are specified, the background copy behavior is the default.

Note: You cannot use the **-nocp** parameter with the **-cp** parameter in the same command.

-seqnum *flash_sequence_num*
(Optional) This parameter is not supported for machine type 2105. Associates the FlashCopy relationships that are established with the specified sequence number. You can use this sequence number as an identifier for a relationship or group of relationships. An example sequence number is `00010`.

-pmir no | required | preferred
Specifies the Remote Pair Copy option that you want to use. The Remote Pair Copy option, sometimes called preserve mirror, preserves synchronous Metro Mirror pairs when the FlashCopy source and target volumes are both Metro Mirror primary volumes and both Metro Mirror secondary volumes on the same storage system. The FlashCopy operation is completed on both the local site and the remote site.

no FlashCopy operations are not completed on the remote site. If the target volume is a Metro Mirror primary volume, the remote copy might temporarily change to the duplex pending state. **no** is the default if the **-pmir** parameter is not specified.

required
FlashCopy operations do not change the state of the Metro Mirror primary volume pair to duplex pending. Both the source Metro Mirror volume pair and the target Metro Mirror volume pair must be in the full duplex state.

preferred
Uses the Remote Pair Copy option for FlashCopy operations when possible. The Remote Pair Copy option cannot be used if the configuration is not correct or the state of the volume is not supported by this function.

-resetreserve
(Optional) Forcibly clears any SCSI reservation on the target volume and allows establishing of a FlashCopy relationship. The reservation is not restored after the relationship is established.

- When this option is not specified and the target volume is reserved, the command fails.
- This option is ignored if the target is a CKD volume; this option is applicable only for fixed block volumes.

source_volume_ID:target_volume_ID ... | -
(Required) Increments a FlashCopy relationship for the source and target volume pairs with the specified IDs. This parameter accepts fully qualified volume IDs, which include storage image IDs, or a shortened version without storage image IDs if the **-dev** parameter is specified or you specify a value for the *devid* variable in your profile file. You must separate multiple FlashCopy pair IDs with spaces.

A FlashCopy pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a FlashCopy relationship. You must separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

X Specifies the address group, 0–F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified FlashCopy pair ID is `IBM.2107-75FA120/0001:IBM.2107-75FA120/0004`

An example of a shortened version is `0001:0004`

An example of multiple pairs is `0001:0004 0003:00FF 0008:000C`

Example

An invocation example

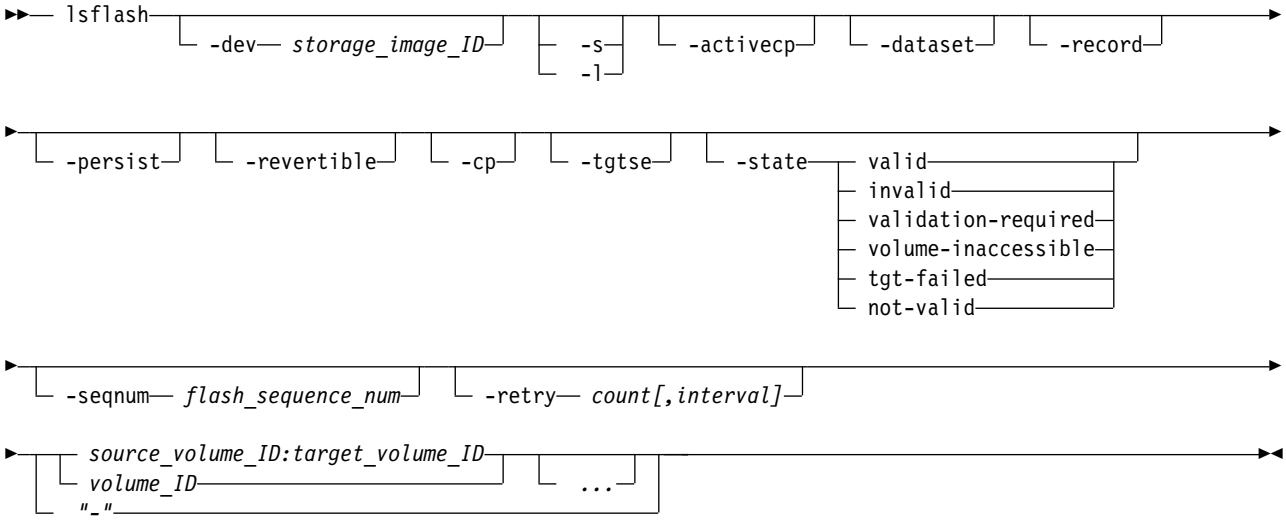
```
dscli> resyncflash –dev IBM.2107-75FA120 -freeze 0100:0200
```

The resulting output

```
Remote FlashCopy volume pair 0100:0200 successfully resynchronized.
```

lsflash

The **lsflash** command displays a list of FlashCopy relationships and status information for each FlashCopy relationship in the list.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the source and target volumes, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-s
(Optional) Displays FlashCopy pair IDs. You cannot use the **-s** and the **-l** parameters together.

-l
(Optional) Displays the default output plus additional attributes that are identified as long output. You cannot use the **-s** and the **-l** parameters together.

-activecp
(Optional) Displays the FlashCopy relationships where their background copy process is active.

Note: The background copy process might be inactive for a while before it starts.

-dataset
(Optional) Displays the volumes that are participating in the Dataset FlashCopy relationships.

-record
(Optional) Displays only the FlashCopy relationships that were established with the –**record** or **-multinc** option.

-persist
(Optional) Displays the FlashCopy relationships that were established with the **-persist** parameter.

-revertible
(Optional) Displays the FlashCopy relationships that were established with the **-revertible** parameter.

-cp
(Optional) Displays the FlashCopy relationships that were established with the **-cp** parameter.

-tgtse
(Optional) Displays the FlashCopy relationships that have a space-efficient target.

-state valid | invalid | validation-required |volume-inaccessible | tgt-failed | not-valid
(Optional) Displays the FlashCopy relationships that are identified by the specific state.

Note: When you specify *not-valid*, all FlashCopy relationships that do not meet the requirements for the *valid* state are displayed.

-seqnum *flash_sequence_number*
(Optional) Displays the FlashCopy relationships that are associated with the specified sequence number. The default is 0000.

Note: This parameter is not supported for machine type 2105.

-retry *count[,interval]*
(Optional) Specifies how you want the system to handle a validation-required state.

The system currently handles a validation-required state as follows:

- If there are one or more FlashCopy relationships, an immediate retry is initiated. In most cases, the reasons for the validation-required state are cleared by the time that the retry is processed and normal processing continues.
- If the validation-required state still exists after the first retry, the system initiates five wait and retry cycles with a delay of 5 seconds between each cycle. At any time during these cycles, if the reasons for the validation-required state are cleared, normal processing continues.

You can change how the system handles a validation-required state as follows:

- Set the number of retries (count) to 0. When you set the number of retries to 0, it prevents the system from attempting any retries.
- Set the number of retries to 1. The system performs an immediate retry if there are one or more FlashCopy relationships in the validation-required state. The 5-second delay is not initiated.
- Set the number of retries to *N*, with *N* greater than 1. The system performs an immediate retry if there are one or more FlashCopy relationships in the validation-required state, followed by at least one wait and retry loop. The default for *N* is 6. You can change the length of the 5-second default wait delay using the optional interval value.

source_volume_ID:target_volume_ID | *volume_ID ...* | **-**
(Required) Displays the FlashCopy relationships for the source and target volume pairs with the specified IDs, or displays the FlashCopy relationships for a single volume ID if the volume ID is specified.

This parameter accepts fully qualified volume IDs, which includes the storage image IDs, or a shortened version without storage image IDs, if the **-dev** parameter is specified or you can specify a value for the *devid* variable in your profile file.

A FlashCopy pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a FlashCopy relationship. You must separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

X Specifies the address group, 0–F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

You must separate multiple IDs with spaces. You can specify FlashCopy pair IDs and a range of FlashCopy pair IDs, or you can specify volume IDs and a range of volume IDs. You cannot specify a combination of FlashCopy pair IDs and volumes IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lsflash** command using the **-l** parameter.

Invoking the lsflash command

```
dscli> lsflash
-dev IBM.2107-75FA120 -l 0100:0200 0101:0201 0102:0202 0103:0203
```

The resulting output

ID	SrcLSS	SequenceNum	Timeout	Active Copy	Recording	Persistent	Revertible
0100:0200	01	10	60	Disabled	Disabled	Disabled	Disabled
0101:0201	01	10	60	Disabled	Disabled	Disabled	Disabled
0102:0202	01	11	60	Disabled	Disabled	Disabled	Disabled
0103:0203	01	11	60	Disabled	Disabled	Disabled	Disabled

Source-Write-Enabled	Target-Write-Enabled	Background-Copy	SGC Recovered	OutOfSyncTracks	DateCreated	DateSynced	State	isTgtSE	Pmir
Enabled	Disabled	Disabled	Yes	0	10/01 /2007 02:20:00	10/01 /2007 02:23:47	Valid	TSE	No
Enabled	Disabled	Disabled	No	0	10/01 /2007 02:20:00	10/01 /2007 02:23:47	Valid	TSE	No
Enabled	Disabled	Disabled	Yes	0	10/01 /2007 02:20:00	10/01 /2007 02:23:47	Valid	ESE	No
Enabled	Disabled	Disabled	No	0	10/01 /2007 02:20:00	10/01 /2007 02:23:47	Valid	ESE	No

Report field definitions

ID Indicates the FlashCopy pair ID. This ID consists of two volume IDs, one designated as the source and the other as the target volume for a FlashCopy relationship. For dataset FlashCopy relationships, the source volume ID is the same as the target volume ID, which indicates that this volume is participating as a source or a target in a dataset FlashCopy relationship.

SrcLSS
Indicates the Consistency Group LSS ID that is associated with the source volume of this FlashCopy relationship.

SequenceNum
Indicates the sequence number that is associated with the FlashCopy relationship.

Note: This item is not supported on the 2105.

Timeout
Indicates the consistency group Long Busy Timeout setting for the LSS ID that is associated with the source volume of this FlashCopy relationship. This value can be modified using the **chlss** (FB) or the **chlcu** (CKD) command.

ActiveCopy
Indicates whether the background copy process is active for this FlashCopy relationship.

Recording
Indicates whether this FlashCopy relationship was created with one of the change recording options. One of the following values is displayed for each FlashCopy relationship:

Table 10. Change recording options

Option	Description
Disabled	Indicates that the relationship was created without any of the change recording parameters.
Enabled	Indicates that the relationship was created without any of the change recording parameters.
MultInc	Indicates a type 1 relationship that was created with the –**record** parameter but without the -**multinc** parameter.

Persistent
Indicates whether this FlashCopy relationship was established with the persistent option.

Revertible
Indicates whether this FlashCopy relationship was established with the revertible option.

SourceWriteEnabled
Indicates whether this FlashCopy relationship was established with the allow source writes option.

TargetWriteEnabled
Indicates whether this FlashCopy relationship was established with the allow target writes option.

BackgroundCopy
Indicates whether this FlashCopy relationship was established with the run background copy option.

SGCRecovered
Indicates whether or not this is an SGC Recovery session. No indicates a flash session.

OutOfSyncTracks
Indicates the number of tracks that are not synchronized for this FlashCopy relationship. The maximum value that can be displayed is dependent on the source volume size. A dash (-) is displayed when the track counter is not available.

DateCreated
Indicates the date and the time that the FlashCopy relationship was established.

DateSynced
Indicates the date and time that FlashCopy relationship was synchronized, or specifies " - " if the relationship is not synchronized.

State Indicates the state of the FlashCopy relationships. One of the following values is displayed for each FlashCopy relationship:

Note: When a query indicates any state other than valid, the only information that is displayed on the report is the FlashCopy pair ID and the state condition. The rest of the information columns are displayed with a " - " value.

Valid Indicates that the FlashCopy relationship is in a normal state, and that it has been queried successfully.

Validation Required
Indicates that the FlashCopy relationship cannot be queried. The reason that the query is blocked is only temporary. If you issue a new query within several seconds, the problem no longer exists.

Tgt Failed
Indicates that the FlashCopy relationship is in an error state. The point-in-time copy is lost, and the FlashCopy relationship must be withdrawn. You must issue the **rmflash** command to remove the FlashCopy relationship.

Volume Inaccessible
Indicates that the volume cannot be accessed and that the query has failed. When this state is displayed, it generally means that the volume is in a fenced condition.

Invalid
Indicates that a general internal error occurred when the query was processed.

Dataset
Indicates that the source volume is participating as a source or a target in a dataset, or extent level, FlashCopy relationship.

Note: For dataset FlashCopy relationships, a dash (-) will be listed for all of the fields except the **ID** and the **State**.

isTgtSE
Indicates whether this FlashCopy relationship has a space-efficient target.

No Indicates that the target is not space-efficient.

TSE Indicates that the target is a track space-efficient volume.

ESE Indicates that the target is an extent space-efficient (ESE) volume. ESE volumes are used for thin provisioning.

Unknown
Indicates that the space allocation method of the target is not known.

Pmir Remote Pair Copy option preserves synchronous Metro Mirror pairs when the FlashCopy source volume and target volume are Metro Mirror primary volumes and the Metro Mirror secondary volumes are on the same storage system. The FlashCopy operation is performed on both the local site and the remote site.

No Indicates that the Remote Pair Copy option was not specified.

Preferred
Indicates that the Remote Pair Copy option was specified. If the target is a Metro Mirror primary, then the FlashCopy function can preserve the Full Duplex mode of the target Metro Mirror relationship, if it is possible. If the Remote Pair Copy function is not possible, you can use processing defined for the Remote Pair Copy option of "No".

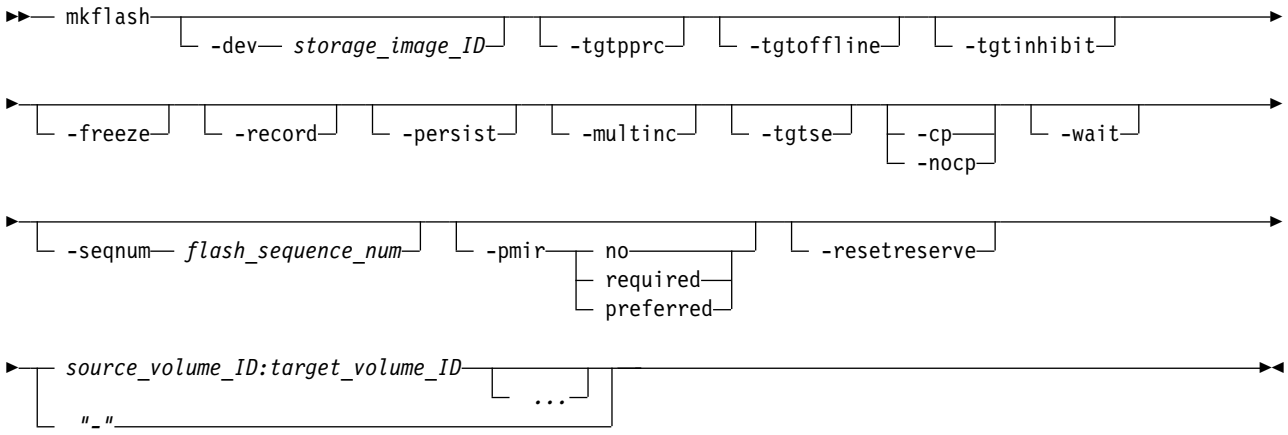
Required
Indicates that the Remote Pair Copy option was specified. If the target is a Metro Mirror primary, then the FlashCopy function is required to preserve the Full Duplex mode of the target Metro Mirror relationship. Processing can fail if the Remote Pair Copy function is not possible.

Remote
Indicates a remote FlashCopy relationship that was initiated by another FlashCopy established at the Metro Mirror primary site with a Remote Pair Copy option of ‘preferred' or ‘required'.

Unknown
The Remote Pair Copy relationship type cannot be determined. The source and target were created with Remote Pair Copy, but they are no longer both Metro Mirror primaries or both Metro Mirror secondaries.

mkflash

The **mkflash** command starts a point-in-time copy from source volumes to target volumes.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the source and target volumes, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command. For example, DS8000 models use IBM.2107-75FA120

-tgtpprc
(Optional) Allows the FlashCopy target volume to be a remote mirror and copy source volume.

-tgtoffline
(Optional) Rejects the **mkflash** command if the target volume is online for host system access. This parameter applies only to count key data (CKD) volumes.

-tgtinhibit
(Optional) Prevents host system write operations to the target volume while the FlashCopy relationship exists.

-freeze
(Optional) Starts the *queue full* condition for the source volume. During the *queue full* condition, the source volume reports a status of *long busy*. All writes to the source volume are queued by the host and are written after the *queue full* condition is reset. The *queue full* condition is reset by an *extended long busy timeout* condition. The timeout condition affects all FlashCopy source volumes that are contained within a respective logical subsystem and that are established or modified with this parameter.

Note: Use the **chlss** and the **chlcu** commands to modify the extended long busy timeout setting.

-record
(Optional) This parameter, without the **-multinc** parameter, creates a type 1 incremental FlashCopy relationship. The type 1 FlashCopy records data changes on both the source and target volumes of the FlashCopy pair.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

Select this parameter when you create an initial FlashCopy relationship that you later want to use with the **resyncflash** or **reverseflash** command. If the **-multinc** parameter is not selected, you can also use the **setflashrevertible** command.

When you select the **-record** parameter, the **-persist** parameter is automatically selected.

-persist
(Optional) Creates a persistent FlashCopy relationship in which the relationship remains after the copy completes and remains indefinitely until a **rmflash** command is issued against the FlashCopy pair. If this parameter is not specified, a normal FlashCopy relationship is created and is automatically removed after the copy completes.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

See the **-record** parameter for a description of a type 1 incremental FlashCopy relationship and the **-multinc** parameter for a description of a type 2 incremental relationship.

When you select either the **-record** or the **-multinc** parameter, the **persist** parameter is automatically selected.

-multinc
(Optional) Creates a type 2 incremental FlashCopy relationship. The type 2 FlashCopy records data changes only on the target volume of the FlashCopy pair.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

The type 2 FlashCopy allows for more than one incremental FlashCopy relationship from the same source volume. However, because the change recording is maintained only on the target volume, the type 2 FlashCopy can cause a performance impact as more type 2 FlashCopy relationships are added.

Select this parameter when you create multiple FlashCopy volume pairs with the same source volume that you want to use with the **resyncflash** and **reverseflash** commands. However, FlashCopy pairs established with this modified recording method cannot be used with the **setflashrevertible** command.

When you select the **-multinc** parameter, the **-persist** and **-record** parameters are automatically selected.

-tgtse
(Optional) Specifies that the target volume that you are designating for a FlashCopy relationship might be a space-efficient logical volume. An error message is generated if the target volume that you are using to create the FlashCopy relationship is a space-efficient volume and you do not specify this parameter.

-nocp
(Optional) Inhibits a background copy. Data is copied from the source volume to the target volume only if a track on the source volume is modified. The FlashCopy volume pair relationship remains indefinitely until it is broken by a **rmflash** command or until all tracks on the source volume are modified.

When the **-tgtse** parameter is specified and neither the **-cp** nor the **-nocp** parameters are specified, the no background copy behavior is the default.

Note: You cannot use the **-nocp** parameter with the **-cp** parameter in the same command.

-cp
(Optional) Starts a background copy. When the **-tgtse** parameter is not specified and neither the **-cp** nor the **-nocp** parameters are specified, the background copy behavior is the default.

Note: You cannot use the **-nocp** parameter with the **-cp** parameter in the same command.

-wait
(Optional) Delays the command response until the background copy process is complete.

Notes:

1. You cannot use the **-wait** parameter with either the **-persist** or the **-nocp** parameters.
2. You cannot use the **-wait** parameter when **-tgtse** is specified and neither the **-cp** nor the **-nocp** parameters are specified. No background copy behavior is the default.

-seqnum *flash_sequence_num*
(Optional) Associates the FlashCopy relationships that are established with the specified sequence number. You can use this sequence number as an identifier for a relationship or group of relationships. An example sequence number is `00010`.

-pmir no | required | preferred
Specifies the Remote Pair Copy option that you want to use. The Remote Pair Copy option, sometimes called preserve mirror, preserves synchronous Metro Mirror pairs when the FlashCopy source volume and target volume are Metro Mirror primary volumes and the Metro Mirror secondary volumes are on the same storage system. The FlashCopy operation is performed on both the local site and the remote site.

no FlashCopy operations are not performed on the remote site. If the target volume is a Metro Mirror primary volume, the remote copy might temporarily change to the duplex pending state. The default is *no* if the **-pmir** parameter is not specified.

required
FlashCopy operations do not change the state of the Metro Mirror primary volume pair to duplex pending. Both the source Metro Mirror volume pair and the target Metro Mirror volume pair must be in the full duplex state.

preferred
Uses the Remote Pair Copy option for FlashCopy operations when possible. The Remote Pair Copy option cannot be used if the configuration is not correct or the state of the volume is not supported with this function.

-resetreserve
(Optional) Forcibly clears any SCSI reservation on the target volume and allows establishing of a FlashCopy relationship. The reservation is not restored after the relationship is established.

- When this option is not specified and the target volume is reserved, the command fails.
- This option is ignored if the target is a CKD volume; this option is applicable only for fixed block volumes.

source_volume_ID:target_volume_ID ... | **-**
(Required) Establishes a FlashCopy relationship for the source and target volume pairs with the specified IDs. This parameter accepts fully qualified volume IDs, which consist of storage image IDs or a shortened version without storage image IDs, if the **-dev** parameter is specified. You can also specify a value for the *devid* variable in your profile file. You must separate multiple FlashCopy pair IDs with spaces.

A FlashCopy pair ID consists of two volume IDs: one designated as the source and the other as the target volume for a FlashCopy relationship. You must separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

Note: You might receive an error message that indicates the number of relationships has been exceeded or that an initial volume format is in progress. This means that the FlashCopy relationship cannot be established because the maximum number of relationships have already been established. Or, the volume was recently created and is still being initialized to support FlashCopy processing. You can issue the **mkflash** command to establish the FlashCopy relationship after the initial volume format process is complete.

The volume ID is a 32-bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ*, where:

X Specifies the address group, 0 - F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

An example of a fully qualified FlashCopy pair ID is `IBM.2107-75FA120/0001:IBM.2107-75FA120/0004`

An example of a shortened version is `0001:0004`

An example of multiple pairs is `0001:0004 0003:00FF 0008:000C`

Example

An invocation example

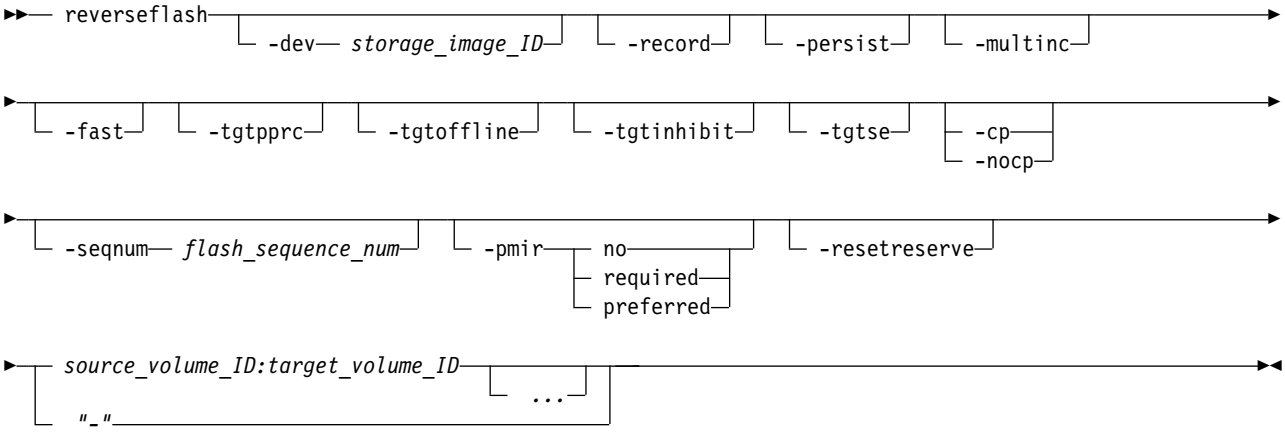
```
dscli> mkflash -dev IBM.2107-75FA120 -freeze 0100:0200
```

The resulting output

```
FlashCopy pair 0100:0200 successfully created.
```

reverseflash

The **reverseflash** command reverses the FlashCopy relationship.



Parameters

You can reverse the direction of a FlashCopy relationship, where the volume that was previously defined as the target becomes the source for the volume that was previously defined as the source. The data that has changed is copied to the volume that was previously defined as the source. For example, you create a

FlashCopy relationship between source volume A and target volume B. Data loss occurs on source volume A. To keep applications running, you can reverse the FlashCopy relationship so that volume B is copied to volume A.

After the reversal takes place, ensure that you designate this new relationship when you issue any future commands. Failure to designate this reversed relationship can produce unexpected results. For example, you use the **reverseflash** command to reverse the relationship of source volume 1600 and target volume 1800. The source volume then becomes 1800 and the target volume becomes 1600. All queries and future processing on this relationship must show volume 1800 as the source and volume 1600 as the target.

The following list defines the parameters that are associated with the **reverseflash** command:

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the source and target volumes, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command. For example, DS8000 models use IBM.2107-75FA120

-record
(Optional) This parameter, without the **-multinc** parameter, creates a type 1 incremental FlashCopy relationship. The type 1 FlashCopy records data changes on both the source and target volumes of the FlashCopy pair.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

Select this parameter when you create an initial FlashCopy relationship that you later want to use with the **resyncflash** or **reverseflash** command. If the **-multinc** parameter is not selected, you can also use the **setflashrevertible** command.

When you select the -**record** parameter, the –**persist** parameter is automatically selected.

-persist
(Optional) Creates a persistent FlashCopy relationship in which the relationship remains after the copy completes and remains indefinitely until a **rmflash** command is issued against the FlashCopy pair. If this parameter is not specified, a normal FlashCopy relationship is created and is automatically removed after the copy completes.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

See the **-record** parameter for a description of a type 1 incremental FlashCopy relationship and the **-multinc** parameter for a description of a type 2 incremental relationship.

When you select either the **-record** or the **-multinc** parameter, the **persist** parameter is automatically selected.

-multinc
(Optional) Creates a type 2 incremental FlashCopy relationship. The type 2 FlashCopy records data changes only on the target volume of the FlashCopy pair.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

The type 2 FlashCopy allows for more than one incremental FlashCopy relationship from the same source volume. However, because the change recording is maintained only on the target volume, the type 2 FlashCopy can cause a performance impact as more type 2 FlashCopy relationships are added.

Select this parameter when you create multiple FlashCopy volume pairs with the same source volume that you want to use with the **resyncflash** and **reverseflash** commands. However, FlashCopy pairs established with this modified recording method cannot be used with the **setflashrevertible** command.

When you select the **-multinc** parameter, the **-persist** and **-record** parameters are automatically selected.

-fast
(Optional) Specify this parameter when you want to issue the **reverseflash** command before the background copy completes.

Note: To use the fast reverse function, the relationship must be set to Target write inhibit. The fast reverse processing function is intended for use as part of Global Mirror recovery process.

At the end of this operation, the original FlashCopy target volume is not usable. Normally, after this command completes the background copy, the new FlashCopy target volume is used as the FlashCopy source volume to restore the original FlashCopy target volume.

-tgtpprc
(Optional) Allows the FlashCopy target volume to be a Remote Mirror and Copy source volume.

-tgtoffline
(Optional) Rejects the **reverseflash** command if the target volume is online for host system access. This parameter applies only to count key data (CKD) volumes.

-tgtinhibit
(Optional) Prevents host system write operations to the target while the FlashCopy relationship exists.

-tgtse
(Optional) Specifies that the target volume that you are designating for a FlashCopy relationship might be a space-efficient logical volume. An error message is generated if the target volume that you are using to create the FlashCopy relationship is a space-efficient volume and you do not specify this parameter.

-nocp
(Optional) Inhibits a background copy. Data is copied from the source volume to the target volume only if a track on the source volume is modified. The FlashCopy volume pair relationship remains indefinitely until it is broken by a **rmflash** command or until all tracks on the source volume are modified.

If neither the **-cp** nor the **-nocp** parameter are specified, then the default background copy behavior is determined by the **-tgtse** and **-fast** parameters. If the **-tgtse** parameter is specified, and the **-fast** parameter is not specified, then the default behavior is to not perform a background copy. Otherwise, the default behavior is to perform a background copy.

Note: You cannot use the **-nocp** parameter with the **-cp** parameter in the same command.

-cp
(Optional) Starts a background copy. Data is copied from the source volume to the target volume as a background task and does not require any source volume modifications to trigger the copy.

If neither the **-cp** nor the **-nocp** parameter are specified, then the default background copy behavior is determined by the **-tgtse** and **-fast** parameters. If the **-tgtse** parameter is specified, and the **-fast** parameter is not specified, then the default behavior is to not perform a background copy. Otherwise, the default behavior is to perform a background copy.

Note: You cannot use the **-nocp** parameter with the **-cp** parameter in the same command.

-seqnum *flash_sequence_num*
(Optional)Associates the FlashCopy relationships that are established with the specified sequence

number. You can use this sequence number as an identifier for a relationship or group of relationships. An example sequence number is `00010`.

-pmir no | required | preferred
Specifies the Remote Pair Copy option that you want to use. The Remote Pair Copy option, sometimes called preserve mirror, preserves synchronous Metro Mirror pairs when the FlashCopy source volume and target volume are Metro Mirror primary volumes and the Metro Mirror secondary volumes are on the same storage system. The FlashCopy operation is performed on both the local site and the remote site.

no FlashCopy operations are not performed on the remote site. If the target volume is a Metro Mirror primary volume, the remote copy might temporarily change to the duplex pending state. The default is *no* if the **-pmir** parameter is not specified.

required
FlashCopy operations do not change the state of the Metro Mirror primary volume pair to duplex pending. Both the source Metro Mirror volume pair and the target Metro Mirror volume pair must be in the full duplex state.

preferred
Uses the Remote Pair Copy option for FlashCopy operations when possible. The Remote Pair Copy option cannot be used if the configuration is not correct or the state of the volume is not supported with this function.

-resetreserve
(Optiona) Forcibly clears any SCSI reservation on the target volume and allows establishing of a FlashCopy relationship. The reservation is not restored after the relationship is established.

- When this option is not specified and the target volume is reserved, the command fails.
- This option is ignored if the target is a CKD volume; this option is applicable only for fixed block volumes.

source_volume_ID:target_volume_ID ... | -
(Required) Establishes a FlashCopy relationship for the source and target volume pairs with the specified IDs. This parameter accepts fully qualified volume IDs, which consist of storage image IDs or a shortened version without storage image IDs, if the **-dev** parameter is specified. You can also specify a value for the *devid* variable in your profile file. You must separate multiple FlashCopy pair IDs with spaces.

A FlashCopy pair ID consists of two volume IDs: one designated as the source and the other as the target volume for a FlashCopy relationship. You must separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ*, where:

X Specifies the address group, 0 - F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

An example of a fully qualified FlashCopy pair ID is `IBM.2107-75FA120/0001:IBM.2107-75FA120/0004`

An example of a shortened version is `0001:0004`

An example of multiple pairs is `0001:0004 0003:00FF 0008:000C`

Example

An invocation example

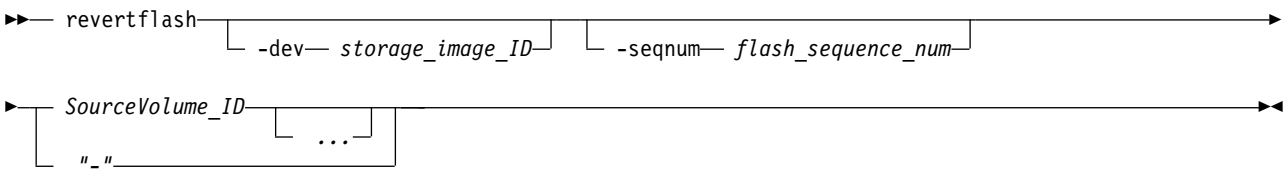
```
dscli> reverseflash -dev IBM.2107-75FA120 0100:0200
```

The resulting output

```
FlashCopy
volume pair 0100:0200 successfully reversed.
```

revertflash

The **revertflash** command is used as part of the recovery from a disaster scenario to rollback a Global Mirror consistency group that is in the process of forming. The former Global Mirror consistency group is restored.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the source volume, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-seqnum *flash_sequence_num*
(Optional) Specifies the FlashCopy sequence number. When this number is specified, the **revertflash** operation is performed only on those relations associated with the specified number.

This parameter is not supported for machine type 2105.

Example: 0010

SourceVolumeID ... | **-**
(Required) Specifies the source volume ID for which the FlashCopy relationship is to be reverted. The chosen FlashCopy pair is the one established or modified with the **-record** parameter. This parameter accepts fully qualified volume IDs, which includes storage image IDs, or a shortened version without storage image IDs if the **-dev** parameter is specified or you specify a value for the *devid* variable in your profile file. You must separate multiple source volume IDs with spaces.

The volume ID is a 32-bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ*, where:

X Specifies the address group, 0 - F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified volume ID: IBM.2107-75FA120/0001

Example of a shortened version: 0001

Example of multiple IDs: 0001 0003 0008

Example

Invoking the revertflash command

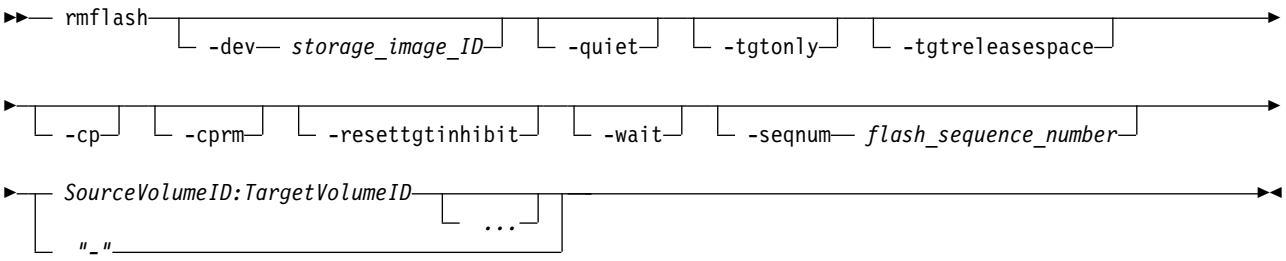
```
dscli> revertflash -dev IBM.2107-75FA120  0100
```

The resulting output

```
FlashCopy pair 0100:0200 successfully reverted.
```

rmflash

The **rmflash** command removes a relationship between FlashCopy volume pairs.



Parameters

Notes:

1. Invoking this command with the **-cp** parameter on a FlashCopy relationship that was previously marked with the **-persist** parameter does not remove the relationship. Instead, the source data is copied to the target.
2. Invoking this command with the **-resettgtinhibit** parameter does not withdraw the relationship, but resets the **-tgtinhibit** parameter if it was previously set.
3. All settings apply to all specified FlashCopy pairs.
4. Do not use the **-wait** parameter on persistent relations.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the source and target volumes, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-quiet
(Optional) Turns off the FlashCopy pair removal confirmation prompt.

-tgtonly
(Optional) Specifies the target volume of the FlashCopy pair to remove the relationship. In addition, the *Copy Indicator* for the target volume is reset.

Note: If you use the **-tgtonly** parameter on CKD volumes, any data set level relationships created by a IBM Z host are removed from the specified target volume.

-tgtreleasespace
(Optional) Specifies that you want the system to release the space that has been allocated, if possible, to a space-efficient logical target volume.

-cp
(Optional) Specifies that the FlashCopy relationship is to be changed from *No Copy* to *Copy* and that the remaining source volume tracks be copied to the target volume. The relationship is removed when all the data is copied unless the relationship is persistent. When this parameter is specified, the copy takes place for all volume pairs where the source volume ID is identical to the source volume that is specified in the command.

-cprm
(Optional) Specifies a change to a FlashCopy relationship from *No Copy* to *Copy*, and copies the remaining source volume tracks to the target volume. The relationship is removed after all of the data is copied. If the FlashCopy relationship is a remote relationship established automatically by another establish FlashCopy with the preserve mirror option set to "preferred" or "required", then the **rmflash** command fails if you do not use the **-cprm** option.

For non-persistent relationships, a Background Copy is initiated for all of the existing relationships within the specified source and target volumes. The relationships are terminated after the background copy is completed. For persistent relationships, all of the existing relationships within the specified source and target volumes change to non-persistent relationships, and then a Background Copy is initiated. The relationships are terminated after the background copy is completed.

-resettgtinhibit
(Optional) Specifies that the parameter that does not allow host system write operations to the target ID while the FlashCopy relationship exists is to be reset, in case it was previously set.

Note: Specifying this parameter in itself does not cause the FlashCopy relationship to be withdrawn.

-wait
(Optional) Specifies that the command response is to be delayed until the background copy process completes.

Notes:

1. Only pairs of source and target volume IDs are allowed when you use the **-wait** parameter.
2. The **-cp** parameter must be used with the **-wait** parameter.
3. Do not use the **-wait** parameter on relationships that are marked **-persist**, an error results from this usage.

-seqnum *flash_sequence_num*
(Optional) Specifies the FlashCopy sequence number. When this number is specified, the **rmflash** operation is performed only on those relationships associated with the specified number.

Example: 0010

SourceVolumeID:TargetVolumeID ... | **-**
(Required) Specifies the source and target volume pairs for which the FlashCopy relationships are removed. This parameter accepts a fully qualified volume ID, which includes storage image IDs, or a shortened version without storage image IDs if the **-dev** parameter is specified or you specify a value for the *devid* variable in your profile file. You must separate multiple FlashCopy pair IDs with spaces.

A FlashCopy pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a FlashCopy relationship. You must separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

When the **-tgtonly** parameter is specified, you must enter volume IDs. Volume pair IDs are not valid with the **-tgtonly** parameter.

The volume ID is a 32 bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

*XY***)** Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use this feature if you are using the DS CLI interactive command mode.

Example of a fully qualified FlashCopy pair ID: IBM.2107-75FA120/0001:IBM.2107-75FA120/0004

Example of a shortened version: 0001:0004

Example of multiple pairs: 0001:0004 0003:00FF 0008:000C

Example

Invoking the rmflash command

```
dscli > rmflash -dev IBM.2107-75FA120 -tgtreleasespace 0100:0200
```

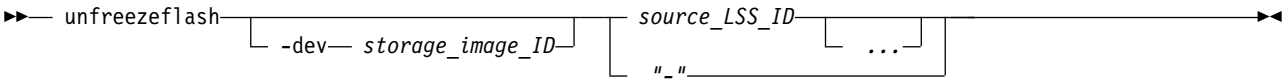
The resulting output

```
Are you sure you want to remove the FlashCopy pair 0001:0002? [y/n]: Y

FlashCopy pair 0100:0200 successfully removed.
```

unfreezeflash

The **unfreezeflash** command resets a FlashCopy consistency group that was previously established with the **-freeze** parameter when the **mkflash** or **resyncflash** commands were issued.



Parameters

-dev *storage_image_ID*

(Optional) Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

source_LSS_ID ... | -

(Required) Specifies that the FlashCopy consistency group be reset for the designated source LSS IDs. The parameter also accepts fully qualified LSS IDs, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified or you specify a value for the *devid* variable in your profile file.

The fully qualified LSS ID format is storage_image_ID/XX. The value for XX is 00 - FE.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input.

Note: You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified LSS ID: IBM.2107-75FA120/00

Example of a shortened version: 00

Example of multiple IDs: 10 20 30

Example

Invoking the unfreezeflash command

```
dscli> unfreezeflash -dev IBM.2107-75FA120 01
```

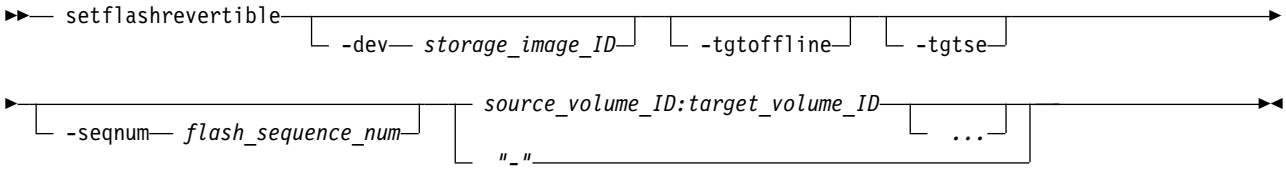
The resulting output

```
FlashCopy consistency group for logical subsystem 01
successfully reset.
```

setflashrevertible

The **setflashrevertible** command modifies a FlashCopy volume pair that is part of a FlashCopy relationship to *revertible*.

The revertible feature allows data to be committed to the target to form a new consistency group or to revert to the last consistency group. This command must be run before the FlashCopy pair can be committed or reverted.



Parameters

Note: The **-nocp**, **-record**, **-persist**, and **-tgtinhibit** (target inhibit) parameters are included automatically when this command processes.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for all source and target volumes, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-tgtoffline
(Optional) Causes an establish FlashCopy volume pair command to be rejected if the target volume is online for host system access. This parameter applies only to CKD volumes.

-tgtse
(Optional) Specifies that the target volume that is part of the FlashCopy relationship that you are modifying to be designated as revertible might be a space-efficient logical volume. An error message is generated if the target volume is a space-efficient volume and you do not specify this parameter.

-seqnum *flash_sequence_num*
(Optional) Associates the FlashCopy relationships that are changed with the specified sequence number. Only the relationships that are successfully modified by the command are assigned the specified sequence number, leaving the ones that fail with the previous number (if previously specified).

Example: 0010

source_volume_ID:target_volume_ID ... | **-**
(Required) Modifies FlashCopy relationships for the source and target volume pairs with the IDs

specified. This parameter accepts fully qualified volume IDs, which includes storage image IDs, or a shortened version without storage image IDs, if the **-dev** parameter is specified. Or, you can specify a value for the *devid* variable that resides in your profile file. You must separate multiple FlashCopy pair IDs with spaces.

A FlashCopy pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a FlashCopy relationship. You must separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32 bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified FlashCopy pair ID:IBM.2107-75FA120/0001:IBM.2107-75FA120/0004

Example of a shortened version: 0001:0004

Example of multiple pairs: 0001:0004 0003:00FF 0008:000C

Example

Invoking the setflashrevertible command

```
dscli> setflashrevertible -dev IBM.2107-75FA120 -tgtse 0100:0200
```

The resulting output

```
FlashCopy volume pair 0100:0200 successfully
made revertible.
```

Remote FlashCopy commands

Commands that are used to configure Remote FlashCopy relationships and to display Remote FlashCopy information are referenced. Remote FlashCopy commands are used to process what was formerly known as inband FlashCopy transactions. These types of transactions cannot be handled through the GUI.

A remote flash is a FlashCopy relationship on a machine at a remote site from the local site machine that you are connected to.

The following Remote FlashCopy commands are available:

commitremoteflash
Sends data to a target volume to form a consistency between the remote source and target FlashCopy pair.

resyncremoteflash
Increments an existing remote FlashCopy pair that has been established with the **-record** and **-persist** parameters.

lsremoteflash
Generates a report that displays a list of FlashCopy relationships and the status information for each FlashCopy relationship in the list.

mkremoteflash
Initiates a remote point-in-time copy from source volumes to target volumes through a remote mirror and copy relationship.

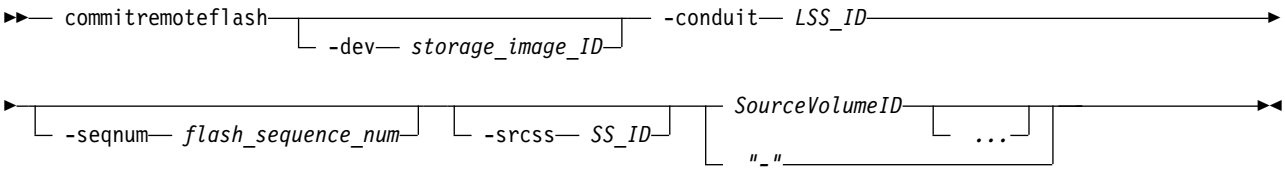
revertremoteflash
Restores data on the source volume to its most recent consistency formation.

rmremoteflash
Removes a relationship between remote FlashCopy volume pairs.

setremoteflashrevertible
Modifies the specified remote FlashCopy volume pair that is part of a Global Mirror relationship to a revertible state. This command must be run before the FlashCopy pair can be committed or reverted.

commitremoteflash

The **commitremoteflash** command sends data to a target volume to form a consistency between the remote source and target FlashCopy pair.



Parameters

Notes:

1. Establish the pair by issuing either the **mkflash** or **mkremoteflash** command with the **-record** and **-persist** parameters.
2. Issue either the **setflashrevertible** or **setremoteflashrevertible** command against the pair.

Only after you have taken these two steps can you issue the **commitremoteflash** command.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-conduit *LSS_ID*
(Required) Specifies the source logical subsystem (LSS) of an existing remote mirror and copy relationship that is to be used as a means for communicating with the remote storage image. The source volume IDs that are specified in the **SourceVolumeID** parameter must serve as secondary volumes in a remote mirror and copy relationship in which one of the conduit LSS volumes serves as a primary volume.

When this parameter is used, you must specify a fully qualified LSS ID. The fully qualified LSS ID format is storage_image_ID/XX. The value for the XX is 00 - FE.

-seqnum *flash_sequence_number*
(Optional) Specifies that the commit operation is performed only on those source volumes that are associated with the specified sequence number.

This parameter is not supported for machine type 2105.

Example: 0010

-srcss *SS_ID*
(Optional) Specifies the subsystem ID of the source logical subsystem at the remote site. When this parameter is used, all source volumes must be within the same logical subsystem.

This parameter is required only for IBM Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: FF10

SourceVolumeID ... | **-**
(Required) Commits remote FlashCopy relationships for the source volumes with the specified IDs. The chosen pair is the one with the enabled **-record** parameter. You must separate multiple source volume IDs with spaces.

This parameter accepts fully qualified volume IDs, which includes the storage image ID, or a shortened version without the storage image ID if either the **-dev** parameter is specified, or you can specify a value for the *devid* variable in your profile file.

The volume ID is a 32 bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified volume ID: IBM.2107-75FA120/0001

Example of a shortened version: 0001

Example of IDs: 0001 0003 0008

Example

Invoking the commitremoteflash command

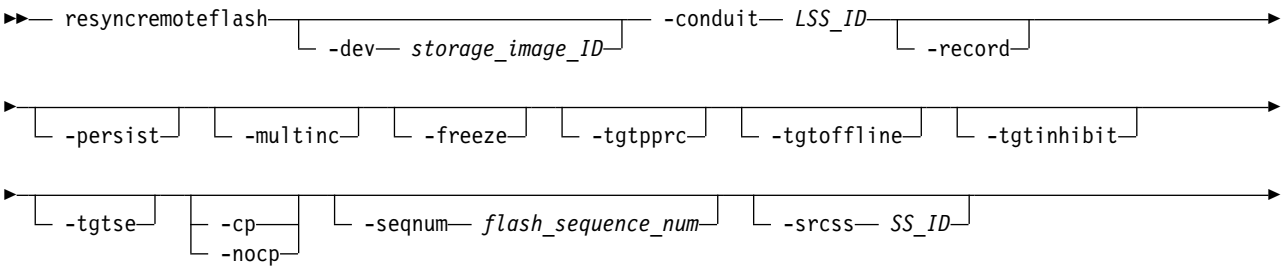
```
dscli> commitremoteflash
-dev IBM.2107-75FA120 -conduit IBM.2107-75FA150/10 0100
```

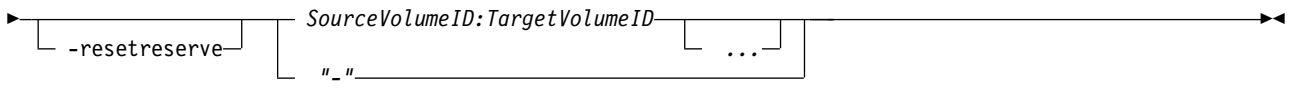
The resulting output

```
FlashCopy pair 0100:0200 successfully committed.
```

resyncremoteflash

The **resyncremoteflash** command (formerly called the **incremoteflash** command and associated with the incremental FlashCopy process) increments an existing remote FlashCopy pair that has been established with the **-record** and **-persist** parameters.





Parameters

Note: When a pair is established with the **-record** and **-persist** parameters, the pair initially synchronizes and then a record of all data that is written from the host to the source is maintained in the source volumes. When the **resyncremoteflash** command is issued on the pair, the new data that is written to the source is copied to the target. The specified parameters in this command replace the parameters in the existing relationship. To keep the initial **-record** and **-persist** parameter values, the **-record** and **-persist** parameters must be specified using the **resyncremoteflash** command.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-conduit *LSS_ID*
(Required) Specifies the source logical subsystem (LSS) of an existing remote mirror and copy relationship that is to be used as a means for communicating with the remote storage image. The source volume IDs that are specified in the **SourceVolumeID:TargetVolumeID** parameter, must serve as secondary volumes in a remote mirror and copy relationship in which one of the conduit LSS volumes serves as a primary volume.

When you use this parameter, you must specify a fully qualified LSS ID. The fully qualified LSS ID format is *storage_image_ID/XX*, where *XX* is a hexadecimal number in the range 00 - FE.

Example: IBM.2107-75FA120/00

-record
(Optional) This parameter, without the **-multinc** parameter, creates a type 1 incremental FlashCopy relationship. The type 1 FlashCopy records data changes on both the source and target volumes of the FlashCopy pair.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

Select this parameter when you create an initial FlashCopy relationship that you later want to use with the **resyncflash** or **reverseflash** command. If the **-multinc** parameter is not selected, you can also use the **setflashrevertible** command.

When you select the -**record** parameter, the –**persist** parameter is automatically selected.

-persist
(Optional) Creates a persistent FlashCopy relationship in which the relationship remains after the copy completes and remains indefinitely until a **rmflash** command is issued against the FlashCopy pair. If this parameter is not specified, a normal FlashCopy relationship is created and is automatically removed after the copy completes.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

See the **-record** parameter for a description of a type 1 incremental FlashCopy relationship and the **-multinc** parameter for a description of a type 2 incremental relationship.

When you select either the **-record** or the **-multinc** parameter, the **persist** parameter is automatically selected.

-multinc
(Optional) Creates a type 2 incremental FlashCopy relationship. The type 2 FlashCopy records data changes only on the target volume of the FlashCopy pair.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

The type 2 FlashCopy allows for more than one incremental FlashCopy relationship from the same source volume. However, because the change recording is maintained only on the target volume, the type 2 FlashCopy can cause a performance impact as more type 2 FlashCopy relationships are added.

Select this parameter when you create multiple FlashCopy volume pairs with the same source volume that you want to use with the **resyncflash** and **reverseflash** commands. However, FlashCopy pairs established with this modified recording method cannot be used with the **setflashrevertible** command.

When you select the **-multinc** parameter, the **-persist** and **-record** parameters are automatically selected.

-freeze
(Optional) Specifies the Freeze Consistency Group condition. This option causes the source volume to be busy (Queue Full status on Open Systems) to all host I/O operations until a FlashCopy Consistency Group Created command is received. All writes to the source volume are queued by the host and are written after the Consistency Group Created command is complete.

During the busy condition, the source volume reports Queue Full for fixed block volumes and busy status for CKD volumes.

The busy condition can also be reset by an *extended long busy timeout* (default 120 seconds). The timeout condition affects all FlashCopy source volumes that are contained within a respective logical subsystem and that are established or modified with the **-freeze** parameter.

Note: This parameter is used with other processing steps for purposes such as backups, testing, or recovery solutions. The use of this parameter ensures that volumes on the target LSSs are consistent with the source LSSs volumes.

-tgtpprc
(Optional) Allows the FlashCopy target volume to be a remote mirror and copy source volume.

-tgtoffline
(Optional) Causes the **resyncremoteflash** command to be rejected if the target volume is online for host system access.

Note: This parameter applies only to count key data volumes.

-tgtinhibit
(Optional) Prevents host system write operations to the target while the FlashCopy relationship exists.

-tgtse
(Optional) Specifies that the target volume might be a space-efficient logical volume. An error message is generated if the target volume is a space-efficient volume and you do not specify this parameter.

-nocp
(Optional) Inhibits background copy. Data is copied from the source volume to the target volume only if a track on the source volume is modified. The FlashCopy volume pair relationship remains indefinitely until it is broken by a **rmflash** command, or until all tracks on the source volume are modified.

When **-tgtse** is specified and the **-cp** and **-nocp** parameters are not specified, the no background copy behavior is the default.

You cannot use the **-nocp** parameter with the **-cp** parameter in the same command.

-cp
(Optional) Specifies that a background copy is to be initiated. When **-tgtse** is not specified and the **-cp** and **-nocp** parameters are not specified, the background copy behavior is the default.

You cannot use the **-cp** parameter with the **-nocp** parameter in the same command.

-seqnum *flash_sequence_num*
(Optional) Associates the FlashCopy relationships that are established with the specified sequence number. You can use this sequence number as an identifier for a relationship or group of relationships. Only the relationships that are modified successfully by the **resyncremoteflash** command are assigned the specified sequence number, leaving the ones that fail with the previous one (if they were previously specified).

This parameter is not supported for machine type 2105.

Example: 0010

-srcss *SS_ID*
(Optional) Specifies the subsystem ID of the source logical subsystem at the remote site. The subsystem ID is a four-digit hexadecimal number in the range (0001 - FFFF). When this parameter is used, all source volumes must be designated within the same logical subsystem.

This parameter is required for IBM Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: FF10

-resetreserve
(Optional) Forcibly clears any SCSI reservation on the target volume and allows establishing of a FlashCopy relationship. The reservation is not restored after the relationship is established.

- When this option is not specified and the target volume is reserved, the command fails.
- This option is ignored if the target is a CKD volume; this option is applicable only for fixed block volumes.

SourceVolumeID:TargetVolumeID ... | -
(Required) Specifies that a remote FlashCopy relationship for the source and target volume pairs be incremented with the designated IDs. This parameter accepts fully qualified volume IDs, which includes storage image IDs or a shortened version without storage image IDs, if the **-dev** parameter is specified,

A FlashCopy pair ID consists of two volume IDs: one designated as the source and the other as the target volume for a FlashCopy relationship. You must separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified FlashCopy pair ID: IBM.2107-75FA120/0001:IBM.2107-75FA120/0004

Example of a shortened version: 0001:0004

Example of multiple pairs: 0001:0004 0003:00FF 0008:000C

Example

Invoking the resyncremoteflash command

```
dscli> resyncremoteflash
-dev IBM.2107-75FA120 -conduit IBM.2107-75FA150/10 0100:0200
```

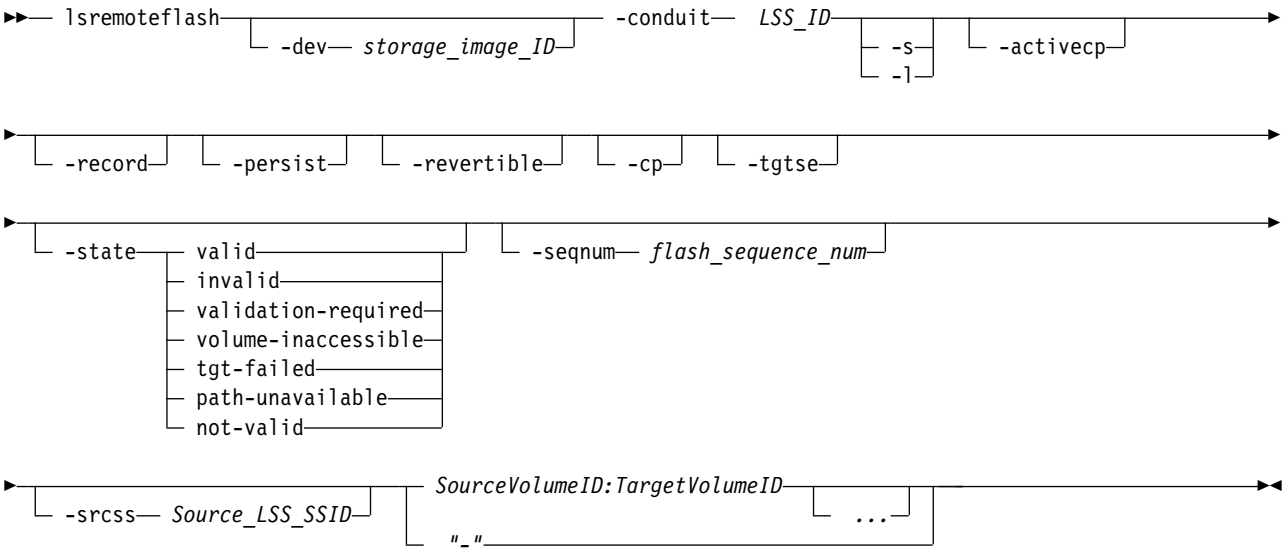
The resulting output

```
Remote FlashCopy volume pair 0100:0200 successfully resynchronized.
 Use the lsremoteflash command to determine copy completion.
```

Note: This message is returned before the copy completes.

lsremoteflash

The **lsremoteflash** command displays a list of remote FlashCopy relationships and status information for each FlashCopy relationship in the list. Remote FlashCopy relationships exist on a remote site, and can be queried using this command from your local site.



Parameters

Note: All settings apply to all FlashCopy pairs that are specified.

-dev *storage_image_ID*
(Optional) Specifies the remote site storage image ID, which consists of manufacturer, machine type, and serial number. The remote site storage image ID is required if you do not specify fully qualified IDs or do not set the *devid* variable in your profile or through the **setenv** command. The remote site storage image ID is also required if the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-conduit *LSS_ID*
(Required) Specifies the source logical subsystem (LSS) of an existing Remote Mirror and Copy relationship that is used as a means for communicating with the remote storage image.

The source volume IDs that are specified in the *SourceVolumeID:TargetVolumeID* parameter must serve as secondary volumes in a Remote Mirror and Copy relationship in which one of the conduit LSS volumes serves as a primary volume.

This parameter accepts a fully qualified LSS ID, which includes the storage image ID. The fully qualified LSS ID format is *storage_image_ID/XX*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

-s
(Optional) Displays only FlashCopy pair IDs. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output plus out-of-sync tracks and date that the FlashCopy relationship was created. You cannot use the **-l** and the **-s** parameters together.

-activecp
(Optional) Specifies that FlashCopy relationships with an active background copy process are to be displayed.

-record
(Optional) Displays only the FlashCopy relationships that were established with the **-record** option.

-persist
(Optional) Specifies that the FlashCopy relationships that were established with the **-persist** parameter are to be displayed.

-revertible
(Optional) Specifies that the FlashCopy relationships that were established with the **-revertible** parameter are to be displayed.

-cp
(Optional) Specifies that the FlashCopy relationships that were established with the run background copy (**-cp**) parameter are to be displayed.

-tgtse
(Optional) Displays the FlashCopy relationships that have a space-efficient target.

-state valid | invalid | validation-required |volume-inaccessible | tgt-failed | path-unavailable | not-valid
(Optional) Displays the FlashCopy relationships that are identified by the specific state.

Note: When you specify *not-valid*, all FlashCopy relationships that do not meet the requirements for the *valid* state are displayed.

-seqnum *flash_sequence_number*
(Optional) Specifies that the FlashCopy relationships that are associated with the specified sequence number are to be displayed.

This parameter is not supported for machine type 2105.

-srcss *Source_LSS_SSID*
(Optional) Specifies the subsystem ID of the source logical subsystem at the remote site. The ID is in the format 0x0001 - 0xFFFF.

This value is required for the IBM Enterprise Storage Server® versions 2.4.0 and 2.4.1. When you specify *SS_IDs*, the source volumes are restricted to one LSS.

Example: FF10

SourceVolumeID:TargetVolumeID ... | **-**
(Required) Specifies that the FlashCopy relationships for the remote site source and target volume pairs with the specified IDs be displayed.

This parameter accepts fully qualified volume IDs, which include storage image IDs, or a shortened version without storage image IDs if the **-dev** parameter is specified.

A FlashCopy pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a FlashCopy relationship. You must separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

You must separate multiple IDs with spaces. You can specify FlashCopy pair IDs and a range of FlashCopy pair IDs, or you can specify volume IDs and a range of volume IDs. You cannot specify a combination of FlashCopy pair IDs and volumes IDs.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified volume ID pair: IBM.2107-75FA120/0001:IBM.2107-68FA120/0004

Example of a shortened version: 0001:0004

Example of multiple pairs: 0001:0004 0003:00FF 0008:000C

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lsremoteflash** command when you use the **-l** parameter.

Invoking the lsremoteflash command

```
dscli> lsremoteflash -l -dev IBM.2107-75FA120
-conduit IBM.2107-75FA150/10 IBM.2107-75FA120/0100:IBM.2107-75FA120/0200
```

The resulting output

ID	SrcLSS	Sequence Num	Active Copy	Recording	Persistent	Revertible
0100:0200	01	10	Disabled	Disabled	Disabled	Disabled

Source-Write-Enabled	Target-Write-Enabled	Back-ground-Copy	Copy-Indicator	OutOf-Sync-Tracks	Date-Created	Date-Synced	State	is-TgtSE	Pmir
Enabled	Disabled	Disabled	Yes	0	10/01 /2007 02:20:00	10/01 /2007 02:23:47	Valid	TSE	No

Report field definitions

ID*
Indicates the FlashCopy pair ID. The FlashCopy pair ID consists of two volume IDs. One is designated as the source and the other is the target volume for a FlashCopy relationship.

SrcLSS
Indicates the logical subsystem ID.

SequenceNum
Indicates the sequence number that is associated with the FlashCopy relationship.

ActiveCopy
Indicates (enabled or disabled) whether the background copy is active on the specified FlashCopy pair.

Recording
Indicates whether this FlashCopy relationship was created with one of the change recording options. One of the following values is displayed for each FlashCopy relationship:

Disabled
Indicates that the relationship was created without any of the change recording parameters.

Enabled
Indicates a type 1 relationship that was created with the **-record** parameter but without the **-multinc** parameter.

MultInc
Indicates a type 2 relationship that was created with both the **-record** and **-multinc** parameters.

Persistent
Indicates (enabled or disabled) whether the designated FlashCopy pair is established with persistent activated.

Revertible
Indicates (enabled or disabled) whether the designated FlashCopy pair is established with the revertible option activated.

SourceWriteEnabled
Indicates (enabled or disabled) whether this FlashCopy relationship was established with the Allow Source Writes option.

TargetWriteEnabled
Indicates (enabled or disabled) whether this FlashCopy relationship was established with the Allow Target Writes option.

BackgroundCopy
Indicates (enabled or disabled) whether this FlashCopy relationship was established with the Run Background Copy option.

OutofSyncTracks+
Indicates the number of tracks that are not synchronized for this FlashCopy relationship. A dash (-) is displayed when the track counter is not available.

DateCreated+
Indicates the date and the time that the FlashCopy relationship was established.

DateSynced
Indicates the date and the time that this FlashCopy relationship was synchronized, or " - " if the relationship is not synchronized.

State
Indicates the state of the FlashCopy relationships. One of the following values is displayed for each FlashCopy relationship:

Note: When a query indicates any state other than valid, the only information that is displayed on the report is the FlashCopy pair ID and the state condition. The rest of the information columns are displayed with a " - " value.

Valid
Indicates that the FlashCopy relationship is in a normal state, and that it has been queried successfully.

Validation Required
Indicates that the FlashCopy relationship cannot be queried. The reason that the query is blocked is only temporary. If you issue a new query within several seconds, the problem no longer exists.

Tgt Failed
Indicates that the FlashCopy relationship is in an error state. The point-in-time copy is lost, and the FlashCopy relationship must be withdrawn. You must issue the **rmflash** command to remove the FlashCopy relationship.

Volume Inaccessible
Indicates that the volume cannot be accessed and that the query has failed. When this state is displayed, it generally means that the volume is in a fenced condition.

Invalid
Indicates that a general internal error has occurred when the query is processed.

Path Unavailable
The specified inband path does not exist. The user must verify that the Remote Mirror and Copy path exists.

isTgtSE
Indicates whether this FlashCopy relationship has a space-efficient target.

No Indicates that the target is not space-efficient.

TSE
Indicates that the target is a track space-efficient volume.

ESE
Indicates that the target is an extent space-efficient (ESE) volume. ESE volumes are used for thin provisioning.

Unknown
Indicates that the space allocation method of the target is not known.

Pmir
The IBM Remote Pair Copy option preserves synchronous Metro Mirror pairs when the FlashCopy source volume and target volume are Metro Mirror primary volumes and the Metro Mirror secondary volumes are on the same storage system. The FlashCopy operation is performed on both the local site and the remote site.

No Indicates that the IBM Remote Pair Copy option was not specified.

Preferred
Indicates that the IBM Remote Pair Copy option was specified. If the target is a Metro Mirror primary, then the FlashCopy function must preserve the Full Duplex mode of the target Metro Mirror relationship, if it is possible. If the IBM Remote Pair Copy function is not possible, you can use processing defined for the IBM Remote Pair Copy option of "No".

Required
Indicates that the IBM Remote Pair Copy option was specified. If the target is a Metro Mirror

primary, then the FlashCopy function is required to preserve the Full Duplex mode of the target Metro Mirror relationship. Processing fails if the IBM Remote Pair Copy function is not possible.

Remote
Indicates that this remote FlashCopy relationship was initiated by another FlashCopy established at the Metro Mirror primary site with the IBM Remote Pair Copy option of 'preferred' or 'required'.

Unknown
The IBM Remote Pair Copy relationship type cannot be determined. The source and target were created with IBM Remote Pair Copy, but they are no longer both Metro Mirror primaries or both Metro Mirror secondaries.

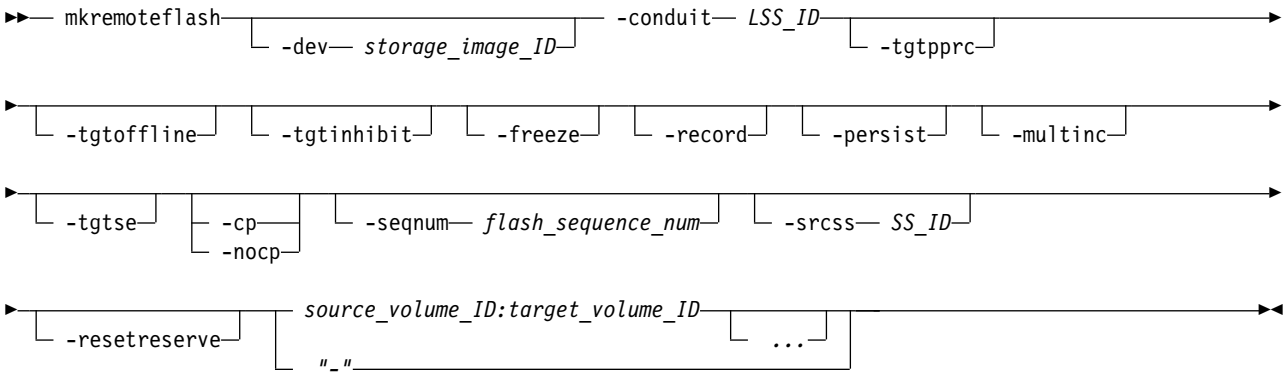
Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

mkremoteflash

The **mkremoteflash** command initiates a remote point-in-time copy from source volumes to target volumes through a Remote Mirror and Copy relationship.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-conduit *LSS_ID*
(Required) Specifies the source logical subsystem (LSS) of an existing remote mirror and copy relationship that is to be used as a conduit for communicating with the remote storage image. The source volume IDs that are specified in the **SourceVolumeID:TargetVolumeID** parameter, must serve as secondary volumes in a remote mirror and copy relationship in which one of the conduit LSS volumes serves as a primary volume.

When you use this parameter, you must specify a fully qualified LSS ID. The fully qualified LSS ID format is storage_image_ID/XX. The value for XX is 00 - FE.

Example: IBM.2107-75FA120/00

-tgtpprc
(Optional) Allows the FlashCopy target volume to be a remote mirror and copy source volume.

-tgtoffline
(Optional) Causes the **mkremoteflash** command to be rejected if the target volume is online for host system access. This parameter applies only to CKD volumes.

-tgtinhibit
(Optional) Prevents host system write operations to the target while the FlashCopy relationship exists.

-freeze
(Optional) Specifies the Freeze Consistency Group condition. The use of this parameter triggers the *queue full* condition for the source volume. All writes to the source volume are queued by the host and are written after the *queue full* condition is reset.

During the *queue full* condition, the source volume reports *long busy* status.

The *queue full* condition is reset by an *extended long busy timeout* condition. The timeout condition affects all FlashCopy source volumes that are contained within a respective logical subsystem and that are established or modified with the **-freeze** parameter.

Note: This parameter is used with other processing steps for purposes such as backups, testing, or recovery solutions. The use of this parameter ensures that volumes on the target LSSs are consistent with the source LSSs volumes.

-record
(Optional) This parameter, without the **-multinc** parameter, creates a type 1 incremental FlashCopy relationship. The type 1 FlashCopy records data changes on both the source and target volumes of the FlashCopy pair.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

Select this parameter when you create an initial FlashCopy relationship that you later want to use with the **resyncflash** or **reverseflash** command. If the **-multinc** parameter is not selected, you can also use the **setflashrevertible** command.

When you select the -**record** parameter, the –**persist** parameter is automatically selected.

-persist
(Optional) Creates a persistent FlashCopy relationship in which the relationship remains after the copy completes and remains indefinitely until a **rmflash** command is issued against the FlashCopy pair. If this parameter is not specified, a normal FlashCopy relationship is created and is automatically removed after the copy completes.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

See the **-record** parameter for a description of a type 1 incremental FlashCopy relationship and the **-multinc** parameter for a description of a type 2 incremental relationship.

When you select either the **-record** or the **-multinc** parameter, the **persist** parameter is automatically selected.

-multinc
(Optional) Creates a type 2 incremental FlashCopy relationship. The type 2 FlashCopy records data changes only on the target volume of the FlashCopy pair.

A single volume can be the source volume to up to 12 FlashCopy relationships, and these relationships can be any combination of normal, persistent, or incremental relationships. However, only a single type 1 incremental relationship can exist.

The type 2 FlashCopy allows for more than one incremental FlashCopy relationship from the same source volume. However, because the change recording is maintained only on the target volume, the type 2 FlashCopy can cause a performance impact as more type 2 FlashCopy relationships are added.

Select this parameter when you create multiple FlashCopy volume pairs with the same source volume that you want to use with the **resyncflash** and **reverseflash** commands. However, FlashCopy pairs established with this modified recording method cannot be used with the **setflashrevertible** command.

When you select the **-multinc** parameter, the **-persist** and **-record** parameters are automatically selected.

-tgtse
(Optional) Specifies that the target volume might be a space-efficient logical volume. An error message is generated if the target volume that you have specified is a space-efficient volume and you do not specify the **-tgtse** parameter.

-nocp
(Optional) Inhibits background copy. Data will be copied from the source volume to the target volume only if a track on the source volume is modified. The FlashCopy volume pair relationship remains indefinitely until it is broken by a **rmremoteflash** command, or until all tracks on the source volume are modified.

When **-tgtse** is specified and the **-nocp** parameter is not specified, the no background copy behavior is the default. You cannot use the **-nocp** parameter with the **-cp** parameter in the same command.

-cp
(Optional) Specifies that a background copy be initiated. When (**-tgtse** is not specified) and neither the **-cp** nor the **-nocp** parameters are specified, the background copy behavior is the default.

You cannot use the **-cp** parameter with the **-nocp** parameter in the same command.

-seqnum *flash_sequence_num*
(Optional) Associates the FlashCopy relationships that are established with the specified sequence number. This sequence number can be used as an identifier for a relationship or group of relationships.

Example: 0010

This parameter is not supported for machine type 2105.

-srcss *SS_ID*
(Optional) Specifies the subsystem ID of the source logical subsystem at the remote site. The ID is in the format 0x0001 - 0xFFFF.

This value is required for the IBM Enterprise Storage Server versions 2.4.0 and 2.4.1. When you specify *SS_IDs*, the source volumes are restricted to one LSS.

Example: FF10

-resetreserve
(Optional) Forcibly clears any SCSI reservation on the target volume and allows establishing of a FlashCopy relationship. The reservation is not restored after the relationship is established.

- When this option is not specified and the target volume is reserved, the command fails.
- This option is ignored if the target is a CKD volume; this option is applicable only for fixed block volumes.

SourceVolumeID:TargetVolumeID ... | -
(Required) Specifies that a remote FlashCopy relationship for the source and target volume pairs be incremented with the designated IDs. This parameter accepts fully qualified volume IDs, which includes the storage image IDs or a shortened version without storage image IDs, if the **-dev** parameter is specified.

A FlashCopy pair ID consists of two volume IDs: one designated as the source and the other as the target volume for a FlashCopy relationship. You must separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified FlashCopy pair ID: IBM.2107-75FA120/0001:IBM.2107-75FA120/0004

Example of a shortened version: 0001:0004

Example of multiple pairs: 0001:0004 0003:00FF 0008:000C

Example

Invoking the mkremoteflash command

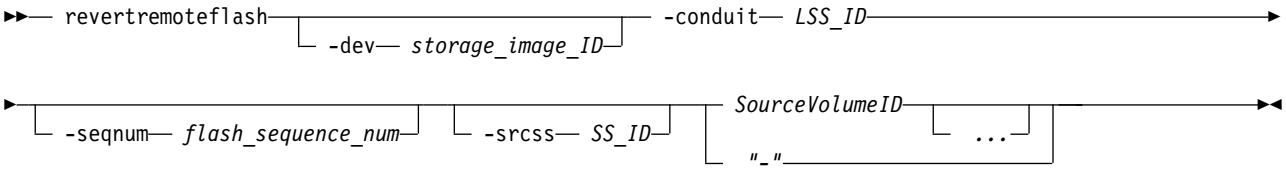
```
dscli> mkremoteflash -dev IBM.2107-75FA120
-conduit IBM.2107-75FA150/10 0100:0200
```

The resulting output

```
FlashCopy volume Pair 0100:0200 successfully created.
Use the lsremoteflash command to determine copy completion.
```

revertremoteflash

The **revertremoteflash** command is used to restore data on the source volume to its most recent consistency formation. All new write operations to the source since the most recent consistency formation are overwritten with the previous consistency.



Parameters

You must take the following actions before you can use the **revertremoteflash** command:

Notes:

1. Issue the **mkflash** or **mkremoteflash** command with the **-persist** and **-record** parameters to establish the FlashCopy pair.
2. Issue the **setflashrevertible** or **setremoteflashrevertible** command against the FlashCopy pair.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs, do not set the

devid variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-conduit *LSS_ID*
(Required) Specifies the source logical subsystem (LSS) of an existing remote mirror and copy (formerly PPRC) relationship that is used as a means for communicating with the remote storage image. The source volume IDs that are specified in *SourceVolumeID:TargetVolumeID* must serve as secondary volumes in a remote mirror and copy relationship in which one of the conduit LSS volumes serves as a primary volume.

When you use this parameter, you must specify a fully qualified LSS ID. The fully qualified LSS ID format is storage_image_ID/XX. The value for XX is 00 - FE.

Example: IBM.2107-75FA120/00

-seqnum *flash_sequence_num*
(Optional) When a FlashCopy sequence number is specified, the **revertremoteflash** operation is performed only on those relationships that are associated with the specified number.

Example: 0010

This parameter is not supported for machine type 2105.

-srcss *SS_ID*
(Optional) Specifies the subsystem ID of the source logical subsystem at the remote site. The ID is in the format 0x0001 - 0xFFFF.

This value is required for the IBM Enterprise Storage Server versions 2.4.0 and 2.4.1. 3. When you specify *SS_IDs*, the source volumes are restricted to one logical subsystem.

Example: FF10

SourceVolumeID ... | -
(Required) Specifies the remote FlashCopy relationship for the source volume with the specified ID that is to be reverted. The chosen FlashCopy pair is the one that is established or modified with the **-record** parameter.

This parameter accepts fully qualified volume IDs, which includes storage image IDs, or a shortened version without storage image IDs if the **-dev** parameter is specified. You must separate multiple source volume IDs with spaces.

The volume ID is a 32-bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a shortened version: 0001

Example of multiple IDs: 0001 0003 0008

Example

Invoking the revertremoteflash command

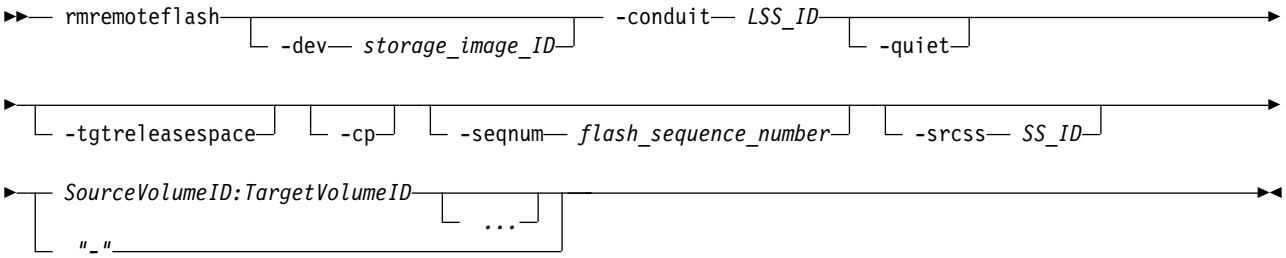
```
dscli> revertremoteflash -dev IBM.2107-75FA120
-conduit IBM.2107-75FA150/10 0100
```

The resulting output

```
Remote FlashCopy volume pair 0100:0200 successfully reverted.
```

rmremoteflash

The **rmremoteflash** command removes a relationship between remote FlashCopy volume pairs.



Parameters

Notes:

1. Invoking this command and using the **-cp** parameter on a FlashCopy relationship that was previously marked with the **-persist** parameter does not remove the relationship. Instead, the source volume is copied to the target volume.
2. Invoking this command resets the **-tgtinhibit** parameter option if it was previously set.
3. All settings apply to all specified FlashCopy pairs.
4. The **-seqnum** parameter is not supported for model 2105.
5. When SS_IDs are specified, the source volumes are restricted to 1 LSS.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-conduit *LSS_ID*
(Required) Specifies the source logical subsystem (LSS) of an existing remote mirror and copy (formerly PPRC) relationship that is to be used as a means for communicating with the remote storage image. The source volume IDs that are specified in *SourceVolumeID:TargetVolumeID* must serve as secondary volumes in a remote mirror and copy relationship in which one of the conduit LSS volumes serves as a primary volume.

This parameter allows the use of a fully qualified LSS ID, which includes the storage image ID. The fully qualified LSS ID format is storage_image_ID/XX. The value for XX is 00 - FE.

-quiet
(Optional) Turns off the FlashCopy pair removal confirmation prompt for this command.

-tgtreleasespace
(Optional) Specifies that you want the system to release the space that has been allocated to the space-efficient target logical volumes back to the repository. This release must occur at the same time that the FlashCopy pair is removed if the only access to the space-efficient volumes is through the conduit LSS ID.

-cp
(Optional) Specifies that the FlashCopy relationship be changed from the *No Copy* to the *Copy* mode. Additionally the remaining source volume tracks are copied to the target volume. The relationship is

removed when all the data is copied unless the relationship is persistent. When the **-cp** parameter is specified, the copy is processed for all volume pairs where the source volume ID is identical to the source volume that is specified in the command.

-seqnum *flash_sequence_num*
(Optional) When a FlashCopy sequence number is specified, the **rmremoteflash** operation is performed only on those relations that are associated with the specified number.

Example: 0010

This parameter is not supported for machine type 2105.

-srcss *SS_ID*
(Optional) Specifies the subsystem ID of the source logical subsystem at the remote site. The ID is in the format 0x0001 - 0xFFFF.

This value is required for the IBM Enterprise Storage Server versions 2.4.0 and 2.4.1. 4. When you specify *SS_IDs*, the source volumes are restricted to one logical subsystem.

Example: FF10

SourceVolumeID:TargetVolumeID ... | -
(Required) Specifies the remote FlashCopy relationships for the source and target volume pairs with the specified IDs that are to be removed.

This parameter accepts fully qualified volume IDs, which includes storage image IDs, or a shortened version without storage image IDs if the **-dev** parameter is specified. You must separate multiple FlashCopy pair IDs with spaces.

A FlashCopy pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a FlashCopy relationship. You must separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified FlashCopy pair ID: IBM.2107-75FA120/0001:IBM.2107-68FA120/0004

Example of a shortened version: 0001:0004

Example of multiple pairs: 0001:0004 0003:00FF 0008:000C

Example

Invoking the rmremoteflash command

```
dscli> rmremoteflash -dev IBM.2107-75FA120
 -conduit IBM.2107-75FA150/10 0100:0200
```

The resulting output

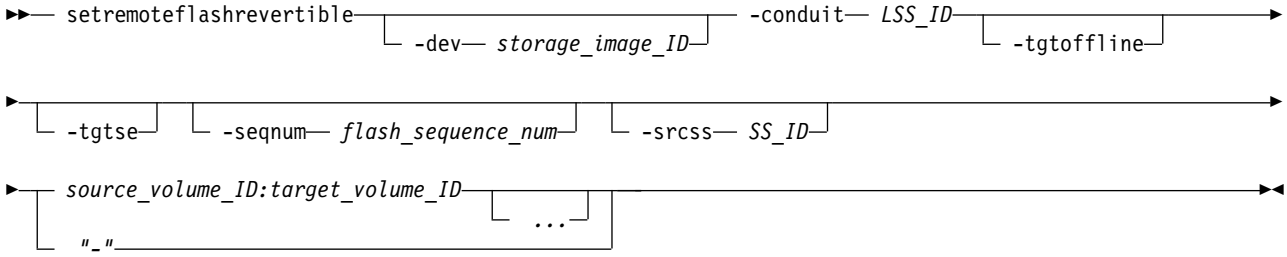
```
Are you sure you want to remove the FlashCopy pair 0100:0200? [y/n]: Y

Removal of the remote FlashCopy volume pair 0100:0200 has been initiated
successfully. Use the lsremoteflash command to determine when the
 relationship is deleted.
```

setremoteflashrevertible

The **setremoteflashrevertible** command modifies a remote FlashCopy volume pair that is part of a FlashCopy relationship to *revertible*.

When a pair is revertible, the data can be committed to the target to form a new consistency group, or it can be reverted back to the last consistency group. This command must be run before the FlashCopy pair can be committed or reverted.



Parameters

Note: The **-nocp**, **-record**, **-persist**, and **-tgtinhibit** (target inhibit) parameters that were specified when the FlashCopy pair was made (**mkremoteflash** command) are included automatically when the **setremoteflashrevertible** command processes. However, FlashCopy pairs established with the modified recording method (multinc) cannot be used with the setflashrevertible command.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-conduit *LSS_ID*
(Required) Specifies the source logical subsystem (LSS) of an existing remote mirror and copy relationship that is to be used as a passage for communicating with the remote storage image. The source volume IDs that are specified in *source_volume_ID:target_volume_ID* must serve as secondary volumes in a remote mirror and copy relationship in which one of the passage LSS volumes serves as a primary volume.

When you use this parameter, you must specify a fully qualified LSS ID. The fully qualified LSS ID format is *storage_image_ID/XX*, where *XX* is a hexadecimal number in the range 00 - FE.

-tgtoffline
(Optional) Causes an establish FlashCopy volume pair command to be rejected if the target volume ID is online for host system access.

This parameter applies only to CKD volumes.

-tgtse
(Optional) Specifies that the target volume that is part of the FlashCopy relationship that you are modifying to be designated as revertible might be a space-efficient logical volume. An error message is generated if the target volume is a space-efficient volume and you do not specify this parameter.

-seqnum *flash_sequence_num*
(Optional) Associates the remote FlashCopy relationships that are changed with the specified sequence number. Only the relationships that are successfully modified by the command get the specified sequence number, leaving the ones that failed with the previous number (if previously specified).

Example: 0010

This parameter is not supported for machine type 2105.

-srcss *SS_ID*
(Optional) Specifies the subsystem ID of the source logical subsystem at the remote site. The ID is in the format 0x0001 - 0xFFFF.

This value is required for the IBM Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: FF10

source_volume_ID:target_volume_ID ... | **-**
(Required) Specifies that the remote FlashCopy relationships for the designated source and target volume pairs be modified.

This parameter accepts fully qualified volume IDs, which includes storage image IDs, or a shortened version without storage image IDs if the **-dev** parameter is specified. You must separate multiple FlashCopy pair IDs with spaces.

A FlashCopy pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a FlashCopy relationship. You must separate the two volume IDs of a FlashCopy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified FlashCopy pair ID: IBM.2107-75FA120/0001:IBM.2107-75FA120/0004

Example of a shortened version: 0001:0004

Example of multiple pairs: 0001:0004 0003:00FF 0008:000C

Example

Invoking the setremoteflashrevertible command

```
dscli> setremoteflashrevertible -dev IBM.2107-75FA120 0100:
```

The resulting output

```
Remote FlashCopy pair 0100:0200 successfully made revertible.
```

Remote Mirror and Copy path commands

Commands are referenced for actions that are related to Remote Mirror and Copy (formerly PPRC).

The following Remote Mirror and Copy path commands are available:

lsavailpprcport
Generates a report that displays a list of ESCON or Fibre Channel I/O ports that can be defined as Remote Mirror and Copy paths.

lspprcpath
Generates a report that displays a list of existing Remote Mirror and Copy path definitions.

mkesconpprcpath
Creates a Remote Mirror and Copy path between source and target logical subsystems over an ESCON connection.

mkpprcpath
Establishes or replaces a Remote Mirror and Copy path between source and target logical subsystems (LSSs) over a Fibre Channel connection.

rmpprcpath
Deletes one or more specified Remote Mirror and Copy paths.

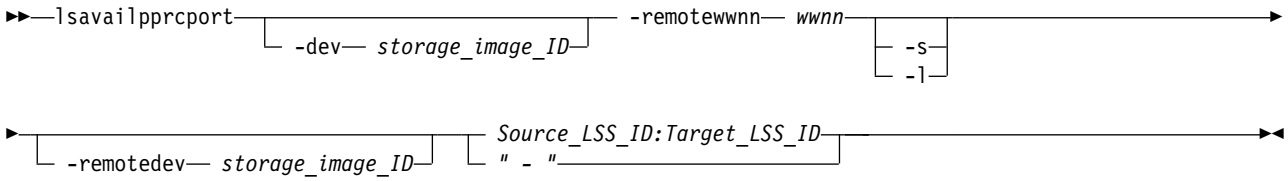
lsavailpprcport

The **lsavailpprcport** command displays a list of ESCON or Fibre Channel I/O ports that can be defined as remote mirror and copy (formerly PPRC) paths.

The DS8000 models support only Fibre Channel ports. The Enterprise Storage Server (2105 machine type) supports ESCON ports.

Secondary ports listed for PPRC are only available when configured as FCP. Secondary FICON ports, while displayed in the list, are not supported, and the DS CLI command mkpprcpath issued to a secondary FICON port will fail.

Note: If you are creating paths between an older release of the DS8000 (Release 5.1 or earlier), which supports only 4-port host adapters, and a newer release of the DS8000 (Release 6.0 or later), which supports 8-port host adapters, the paths should connect only to the lower four ports on the newer storage system.



Parameters

-dev *storage_image_ID*
(Optional). Specifies the source volume storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-remotewwnn *wwnn*
(Required) Specifies the worldwide node name of the secondary storage image. The format is a 16-hexadecimal ID or a colon-separated string.

Example: 12341234000A000F or 12:34:12:34:00:0A:00:0F

Note: You want to use the WWNN that is associated with the remote storage image. Run the **lssi** or **showsi** command to obtain this number. If you use the WWNN that is associated with the primary storage system, this command fails.

-s
(Optional). Displays the local port ID. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional). Displays all fields. You cannot use the **-l** and the **-s** parameters together.

-remotedev *storage_image_ID*
(Required or Optional). Specifies the remote storage system that contains the I/O ports that are queried by the *Source_LSS_ID:Target_LSS_ID* parameter. The remotedev ID consists of the value for the manufacturer, machine type, and serial number.

Required - This parameter is required when querying ESCON I/O ports unless a fully qualified target logical subsystem ID is specified.

Optional - This parameter is optional if you are querying fibre channel I/O ports.

Note: If specified the format of this entry might be checked for correctness even though the value is not used.

Example: IBM.2107-75FA120

Source_LSS_ID:Target_LSS_ID | **-**
(Required) Queries I/O ports that are available for a remote mirror and copy path relationship for the source and target LSSs. This parameter accepts fully qualified LSS IDs, which includes the storage image ID or shortened version without the storage image ID, if the **-dev** parameter is specified.

A remote mirror and copy path LSS pair ID consists of two LSS IDs, one designated as the source LSS and the other as the target LSS for a remote mirror and copy path relationship. The two LSS IDs must be separated with a colon and no spaces. The first LSS ID is the source LSS. The second LSS ID is the target LSS.

The fully qualified LSS ID format is *storage_image_ID/XX*, where *XX* is a hexadecimal number in the range 00 - FE.
If you do not use the **-dev** and **-remotedev** parameters, the fully qualified *source_LSS_ID:target_LSS_ID* value must be placed after the **-remotewwnn** value in your command line. Your command line can look like the following example:

```
dscli> lsavailpprcport –l
–remotewwnn 12341234000A000F IBM.2107-75FA120/01:IBM.2107-75FA150/01
```

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example pair: 00:00

Example of multiple pairs: 00:00 01:01 02:02

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lsavailpprcport** command using the **-l** parameter.

Invoking the lsavailpprcport command

```
dscli> lsavailpprcport –l
–dev IBM.2107-75FA120 –remotewwnn 12341234000A000F 01:01
```

The resulting output

ESCON port information displays for the 2105 machine type.

Local port	Attached port	Type	Switch ID	Switch port
I0100	I0200	FCP	N/A	N/A
I0150	I0620	ESCON	N/A	N/A

Local port	Attached port	Type	Switch ID	Switch port
I0200	N/A	ESCON Switch	IBM.111.2222. 75113AB	I10
I0250	N/A	ESCON Switch	IBM.111.2222. 75113AB	I20

Report field descriptions

Local port
Indicates the fully qualified unique Port ID on the local storage system. FCP and ESCON port IDs are displayed as follows:

FCP port ID
Four hexadecimal characters in the format 0xEEAP, where 'EE' is a port enclosure number (00 - 3F), 'A' is the adapter number (0 - F), and 'P' is the port number (0 - F). The FCP port ID number is prefixed with the letter I.

ESCON port ID
Four hexadecimal characters in the format 0xEEAP, where 'EE' is a port enclosure number (04 - 07), 'A' is the adapter number (0 - 3), and 'P' is the port number (0 - 1). The ESCON port ID number is prefixed with the letter I.

Note: When you specify the **-s** parameter, the local port information is the only information displayed on the report.

Attached port
Indicates the fully qualified unique Port ID on the attached storage system. FCP and ESCON port ID numbers are displayed in the format that is described for Local port. However, if there is an ESCON Switch value, the value displayed in this column is **N/A** (not applicable).

Type Indicates the connection type. FCP is the only applicable value for a 2107 or 1750 machine type. For a 2105 machine type, you can have a value of ESCON or ESCON Switch.

Switch ID
Indicates the Switch ID for ESCON Switch connections.

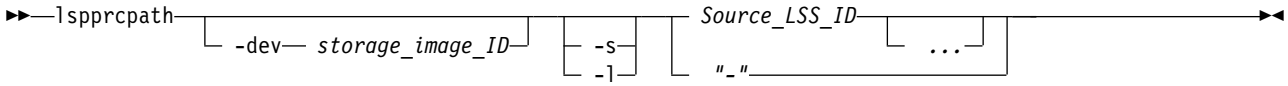
Note: For FCP and direct ESCON, the displayed value in this field is **N/A** (not applicable).

Switch port
Indicates the Port ID on the Switch device that is connected to the attached ESS. The Switch port ID component is two hexadecimal characters in the format 0xPP, where 'PP' is a port number (00 - ff). The number is prefixed with the letter I.

Note: For FCP and direct ESCON, the value of this field is **N/A** (not applicable).

lspprcpath

The **lspprcpath** command displays a list of existing remote mirror and copy path definitions.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of values for the manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID

for the source LSS, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-s
(Optional) Displays the default output of the report but does not include the *Failed Reason* column. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output and the *Failed Reason* and *PPRC CG* descriptions. You cannot use the **-l** and the **-s** parameters together.

Source_LSS_ID ... | -
(Required) Specifies that the Remote Mirror and Copy paths that are defined for the specified source LSS IDs be displayed.

This parameter accepts ranges and individual LSS IDs. You might specify fully qualified LSS IDs, including the storage image ID, or a shortened version if the **-dev** parameter is specified. The fully qualified LSS ID format is *storage_image_ID/XX*, where*XX* is a hexadecimal number in the range 00 - FE.

You must separate multiple LSS IDs with spaces.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example: 00

Example of multiple source LSS IDs: 00 01 02

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report when the **-l** parameter is used with the **lspprcpath** command.

Invoking the lspprcpath command

```
dscli> lspprcpath -dev IBM.2107-75FA120 -l 10
```

The resulting output

Src	Tgt	State	SS
IBM.2107-75FA120/10	IBM.2107-75FA150/10	Failed	0010
IBM.2107-75FA120/10	IBM.2107-75FA150/11	Success	0011
IBM.2107-75FA120/10	IBM.2107-75FA150/12	Degraded	0012
IBM.2107-75FA120/10	IBM.2107-75FA150/13	Invalid	0013

Port	Attached Port	Tgt WWNN	Failed Reason	PPRC CG
IBM.2107-75FA120/I0100	IBM.2107-75FA150/I0100	3007ACF 3012399E0	Configuration Error	Disabled
IBM.2107-75FA120/I0100	IBM.2107-75FA150/I0100	3007ACF-3012399E0	-	Enabled

Port	Attached Port	Tgt WWNN	Failed Reason	PPRC CG
IBM.2107-75FA120/I0100	IBM.2107-75FA150/I0100	3007ACF-3012399E0	Path Degraded Due to High Failure Rate	Disabled
IBM.2107-75FA120/I0100	IBM.2107-75FA150/I0100	3007ACF-3012399E0	-	Enabled

Report field definitions

Src Indicates the fully qualified logical subsystem ID. Use the **-fullid** parameter to display fully qualified IDs, which include the storage image ID.

Tgt Indicates the fully qualified remote mirror and copy ID for the path target LSS. Use the **-fullid** parameter to display fully qualified IDs, which include the storage image ID.

State Displays the current remote mirror and copy path state. One of the following values can be displayed:

Failed The path is not established and has failed. When this is the state value, see the Failed Reason column for an explanation.

Success
The path is established and it is operating normally. When this is the state value, the Failed Reason column displays a " - " value.

Invalid
The path is in an unknown state. When this is the state value, the Failed Reason column displays a " - " value.

Degraded
The path is established, but with degraded performance. When this is the state value, see the Failed Reason column for an explanation.

SS Indicates the subsystem identifier (SSID) of the target LSS.

Port Indicates the fully qualified unique Port ID for the source storage system.

The port ID component is four hexadecimal characters in the format *EEAP*, where *EE* is a port enclosure number (00 - 3F), *A* is the adapter number (0 - F), and *P* is the port number (0 - F). The number is prefixed with the letter *I*.

Attached Port
Indicates the fully qualified unique Port ID for the attached secondary storage system.

The port ID component is four hexadecimal characters in the format 0x*EEAP*, where *EE* is a port enclosure number (00 - 3F), *A* is the adapter number (0 - F), and *P* is the port number (0 - F). The number is prefixed with the letter *I*.

Tgt WWNN
Indicates the worldwide node name of the remote storage image.

Failed Reason
Indicates the reason for the path state. You must issue the **lspprcpath** command with the **-l** parameter to see the values displayed in this field. If the State field has a value of Invalid or Success, a " - " value is displayed in this field. When the State field displays a value of **Failed**, one of the following values is displayed:

Configuration Error
A path has failed for one of the following reasons:

- The specification of the SA ID does not match the installed ESCON adapter cards in the primary controller.
- For ESCON paths, the secondary control unit destination address is zero and an ESCON Director (switch) was found in the path.
- For ESCON paths, the secondary control unit destination address is nonzero and an ESCON Director does not exist in the path. That is, the path is a direct connection.

Delete the original entry and resubmit the **mkpprcpath** command.

Down An FCP path has failed because of a communication or hardware failure.

Primary Login Exceeded
The maximum number of log ins for each source FCP path has been exceeded.

Retry Exceeded
The maximum number of times that the storage system tried to reestablish FCP paths has been exceeded.

Secondary Login Exceeded
The maximum number of log ins for each FCP path to the secondary LSS has been exceeded. The FCP target is unavailable.

Secondary Unavailable
An FCP path to the secondary LSS is unavailable.

Primary No Resources
No resources are available at the source site for the logical paths to be established.

Retry Indicates the number of attempts to reestablish path connection.

Secondary Mismatch
Indicates that there is a mismatch that involves the secondary control unit sequence number or the LSS.

Secondary No Resources
Indicates that resources are not available at the secondary LSS to establish logical paths.

Secondary LSS Mismatch
Indicates that there is a mismatch of the secondary control unit LSS ID or a failure of the I/O that collects secondary information for validation.

Timeout
Indicates that a timeout has occurred. No reason is available.

Not Properly Configured
Indicates that the primary Fibre Channel adapter is not configured properly, or it is not loaded with the correct version of microcode.

Secondary Not PPRC Capable
Indicates that the Fibre Channel path from secondary adapter is not capable of processing a remote mirror and copy path. This can occur from one of the following reasons:

- The secondary adapter is not configured properly, or it is not loaded with the correct version of microcode.
- The secondary adapter is already a target of 32 different storage systems.

ESCON Channel Direction
Indicates that the primary control unit port or link cannot be converted to channel mode because a logical path is already established on the port or link. The establish path operations are not automatically retried within the control unit.

ESCON Initialization Failed
Indicates that initialization for the ESCON protocol has failed.

ESCON Link Offline
Indicates that the ESCON link is offline. This is caused by the lack of light detection coming from a host, peer, or switch.

Path Degraded Due to High Failure Rate
Indicates that a Fibre Channel path is established; however, because of the high failure rate, the path is degraded.

Path Removed Due to High Failure Rate
Indicates that the Fibre Channel path link has been removed because the path has experienced a high failure rate.

System Reserved Path
Indicates that the system has reserved resources for a remote mirror and copy path, for example, after a **failoverpprc** or a **freezepprc** command is used. The resources can be used later, such as for the **failbackpprc** command. In most cases, no action is required. If it is known that a system reserved path is not required, it can be removed with the **rmpprcpath** command only after there are no remote mirror and copy pairs remaining between the LSSs.

PPRC CG
Displays the status of the PPRC consistency group. You must issue the **lspprcpath** command with the **-l** parameter to see the values displayed in this field. One of the following values can be displayed:

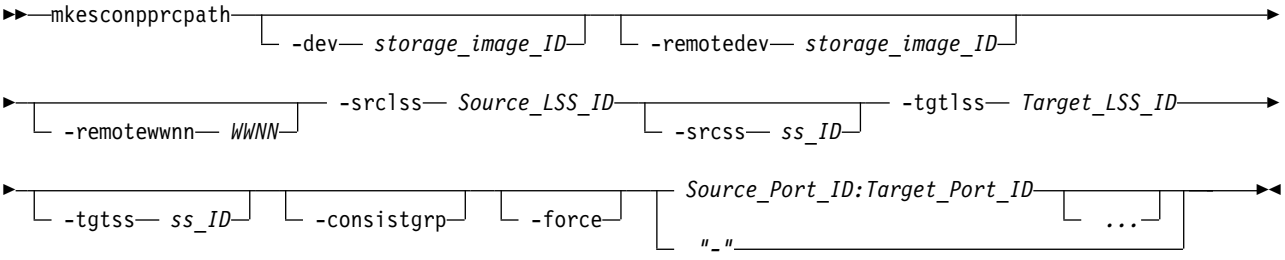
Enabled
Indicates that the remote mirror and copy consistency group is enabled.

Disabled
Indicates that the remote mirror and copy consistency group is disabled.

mkesconpprcpath

The **mkesconpprcpath** command creates a remote mirror and copy (formerly PPRC) path between source and target logical subsystems over an ESCON connection.

The command allows you to specify ESCON direct and ESCON switch connections. Use this command only with IBM Enterprise Storage Servers (2105, Model 800 and Model 750).



Parameters

Notes:

1. The **mkesconpprcpath** command is applicable **only** for the IBM Enterprise Storage Server (2105, Model 800 and Model 750). DS8000 models support only Fibre Channel connections.
2. When you specify a switch port ID as the target port, specify the outgoing port that is connected to the remote ESS and not to the incoming port that is connected to the local ESS.

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified source LSS ID, do

not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2105-FA120

-remotedev *storage_image_ID*
(Optional). Specifies the remote storage image ID, which consists of manufacturer, machine type, and serial number. This parameter is required if you do not fully qualify the target LSS ID.

Example: IBM.2105-FA150

-srclss *Source_LSS_ID*
(Required). Specifies the source logical subsystem (LSS) ID. Accepts a fully qualified LSS ID, which includes the storage image ID or a shortened version without the storage image ID, if the **-dev** parameter is used. The fully qualified LSS ID format is storage_image_ID/*xx*, where '*xx*' is a hexadecimal number in the range '00 - FE'.

-srcss *ss_ID*
(Optional). Specifies the subsystem ID of the primary logical subsystem in the format '0x0001 - 0xFFFF'.

This value is required for the IBM Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-tgtlss *Target_LSS_ID*
(Required). Specifies the target logical subsystem (LSS) ID. Accepts a fully qualified LSS ID, which includes the storage image ID, or a shortened version without the storage image ID, if the **-remotedev** parameter is used. The fully qualified LSS ID format is storage_image_ID/*xx*, where '*xx*' is a hexadecimal number in the range '00 - FE'.

-remotewwnn *WWNN*
(Optional). Specifies the worldwide node name. The format is a 16-hexadecimal ID.

Note: If you use this parameter, the format is checked even though there might be times that the value is not used.

Example: 12341234000A000F

-tgtss *ss_ID*
(Optional). Specifies the subsystem ID of the secondary logical subsystem in the format '0x0001 - 0xFFFF'.

This value is required for the IBM Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-consistgrp
(Optional). Creates a consistency group for the remote mirror and copy volume pairs that are associated with the PPRC paths that are established by this command. A remote mirror and copy consistency group is a set of remote mirror and copy volume pairs that have the same source and target LSS.

Normally, when an error occurs in a member of a remote mirror and copy volume pair, the volume is put in a *suspended* state. However, if the volume is participating in a consistency group, it is placed in a *long busy* state.

-force
(Optional). Creates a new remote mirror and copy path even if the specified remote mirror and copy path already exists.

Source_Port_ID:Target_Port_ID ... | **-**
(Required). Establishes a remote mirror and copy path between the source and target ports for the

specified source and target logical subsystems. The source port must be an ESCON I/O port that is configured for point-to-point or switch topology. The source port is enabled automatically for remote mirror and copy primary I/O operations. The target port must be a switch I/O port that is configured for point-to-point or switch topology. The target port is enabled automatically for remote mirror and copy primary I/O operations.

Note: Do not specify a target port ID when you specify an ESCON direct connection. Instead, specify only the source port ID.

This parameter accepts only non-fully qualified port IDs, which does not include the storage image ID. A remote mirror and copy path port pair ID consists of two port IDs. The first is designated as the source port and the second as the target port for the remote mirror and copy path. You must separate the two port IDs with a colon and no spaces. A direct ESCON I/O port ID is four hexadecimal characters in the format *EEAP*, where *EE* is a port enclosure number '00 - 3F', *A* is the adapter number '0 - F', and *P* is the port number '0 - F'. This number is prefixed with the letter *I*. A switch ESCON I/O port ID is two hexadecimal characters in the range '00 - FF'. This number is prefixed with the letter *I*.

This parameter accepts up to eight remote mirror and copy path port pair IDs. You must separate multiple port pair IDs with spaces.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example pair: I1A10:I20

Example of a source ESCON port and target switch port pair: I1A10:I20

Example of multiple pairs: I1A10:I20 I1A11:I21 I1A12 (the last object identifies an ESCON connection)

Example (2105 use only)

Invoking the mkesconpprcpath command

```
dscli> mkesconpprcpath -dev IBM.2105-FA120
-remotedev IBM.2105-FA150 -srclss 01 -tgtlss 01 I0100:I20 I0110:I21
```

The resulting output

```
Remote Mirror and Copy path
IBM.2105-FA120/01:IBM.2105-FA150/01 successfully created.
```

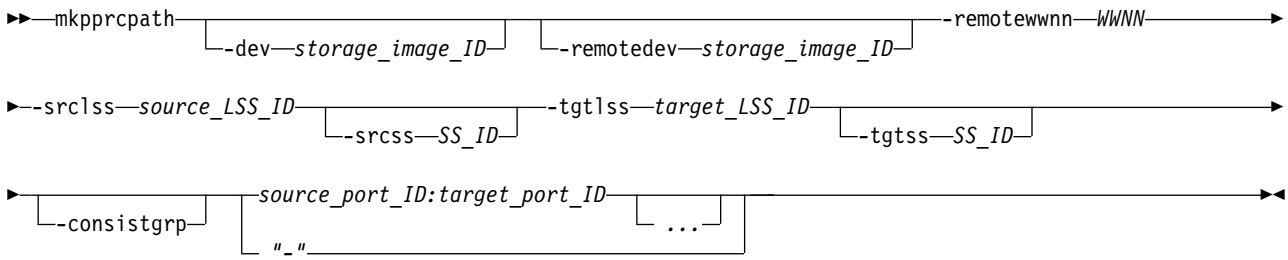
mkpprcpath

The **mkpprcpath** command establishes or replaces a remote mirror and copy (formerly PPRC) path between source and target logical subsystems (LSSs) over a Fibre Channel connection.

It is not necessary to remove any existing remote mirror and copy paths before you establish new paths. All existing paths are automatically removed before the new paths are established. As a result, all multiple paths between any specific pair of source and target LSSs must be specified in a single **mkpprcpath** command.

Machine type 2107 supports this connectivity. Paths can be established between the following machine types: 2105:2105, 2107:2107, 2107:1750, 2107:2105, 1750:1750, 1750:2105.

Note: Secondary ports that are listed for PPRC, after you enter the **lsavailpprcport** command, are only available when configured as FCP. Secondary FICON® ports, while displayed in the list, are not supported, and the DS CLI command **mkpprcpath** entered to a secondary FICON port fails.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs, do not set the *devid* variable in your profile or through the **setenv** command. The storage image ID is also required if the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-remotedev *storage_image_ID*
(Optional) Specifies the ID of the secondary storage image, which includes manufacturer, machine type, and serial number. If specified, the format of this entry might be checked for correctness although the value is not used.

Example IBM.2107-75FA150

-remotewwnn *WWNN*
Required) Specifies the worldwide node name of the secondary storage image. The format is a 16-hexadecimal ID or a colon-separated string.

Note: Ensure that you use the worldwide node name that is associated with the secondary storage system. Run the **lssi** or **showsi** command to obtain this number.

Example 12341234000A000F or 12:34:12:34:00:0A:00:0F

-srclss *source_LSS_ID*
(Required) Specifies the source logical subsystem ID. Use a fully qualified LSS ID, which includes the storage image ID, or use a shortened version without the storage image ID, if the **-dev** parameter is used. The fully qualified LSS ID format is *storage_image_ID/XX*, where *XX* is a hexadecimal number in the range 00 - FE.

Example of a shortened version 00

-srcss *SS_ID*
(Optional) Specifies the subsystem ID of the primary logical subsystem in the format 0x0001 - 0xFFFF.

This value is required for the IBM TotalStorage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example 0010

-tgtlss *target_LSS_ID*
(Required) Specifies the logical subsystem ID associated with the secondary storage system as the target. Use a fully qualified LSS ID, which includes the storage image ID. The fully qualified LSS ID format is *storage_image_ID/XX*, where *XX* is a hexadecimal number in the range 00 - FE.

Example of a shortened version 01

-tgtss *SS_ID*
(Optional) Specifies the subsystem ID of the secondary logical subsystem in the format 0x0001 - 0xFFFF.

This value is required for the IBM TotalStorage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example 0010

-consistgrp
(Optional) Creates a consistency group for the remote mirror and copy volume pairs. A remote mirror and copy consistency group is a set of remote mirror and copy volume pairs that have the same source and target LSS. This option is intended to be used only with Metro Mirror pairs.

Normally, when an error occurs in a member of a remote mirror and copy volume pair, the storage system places the volume in a *suspended* state. However, if the volume participates in a consistency group, it is placed in a *long busy* state.

Note: A consistency group volume pair that is suspended will not go into an extended long busy state unless the primary volume is in a full duplex state before the suspension.

source_port_ID:target_port_ID ... | -
(Required) Establishes a remote mirror and copy path between the source and target ports for the specified source and target logical subsystems. The source and target ports must be Fibre Channel I/O ports that are configured for point-to-point or switched fabric topology. They are enabled automatically for remote mirror and copy secondary I/O operations. They are not enabled for FICON I/O operations.

Use fully qualified port IDs, which include the storage image ID, or use a shortened version without the storage image ID if the **-dev** parameter is specified. A remote mirror and copy path port pair ID consists of two port IDs. Designate the first as the source port and the second as the target port for the remote mirror and copy path. You must separate the two port IDs with a colon and no spaces. A port ID is four hexadecimal characters in the format *EEAP*, where *EE* is a port enclosure number (00 - 3F), *A* is the adapter number (0 - F), and *P* is the port number (0 - F). This number is prefixed with the letter **I**.

To establish multiple remote mirror and copy paths between any single source and target LSS, you must specify all of the path port pair IDs in a single **mkpprcpath** command. This parameter accepts up to eight remote mirror and copy path port pair IDs. You must separate multiple port pair IDs with spaces.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of the shortened version - I1A10:I2A20

Example of multiple pairs - I1A10:I2A20 I1A11:I2A21 I1A12:I2A22

Example

Invoking the mkpprcpath command

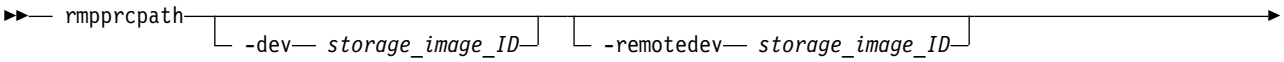
```
dscli> mkpprcpath -dev IBM.2107-75FA120
-srclss 01 -tgtlss 01 -remotewwnn 12341234000A000F I0100:I0100
```

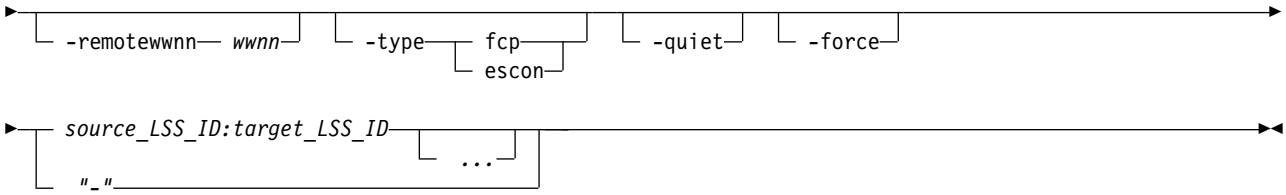
The resulting output

```
Remote Mirror and Copy path 01:01 successfully established.
```

rmpprcpath

The **rmpprcpath** command deletes a Remote Mirror and Copy path.





Parameters

-dev *storage_image_ID*
(Optional) Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified source LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-remotedev *storage_image_ID*
(Optional) Specifies the target storage image ID, which includes manufacturer, machine type, and serial number. This parameter is required if you do not specify a fully qualified target LSS ID or if the **-dev** parameter is used.

-remotewwnn *wwnn*
(Optional) Specifies the secondary worldwide node name. The format is a 16-hexadecimal ID or a colon-separated string.

Note: The following considerations can help you decide whether to use this parameter:

- If you do not specify this parameter, DS CLI processing requires a query for this information from the remote device. In some cases, due to the path-specific state, the query might fail to locate the remote WWNN. If the remote WWNN cannot be located, the **rmpprcpath** command fails. Process the **lspprcpath** command to obtain the remote WWNN information and then process the **rmpprcpath** command with the remote WWNN information included.
- Use the **lspprcpath** command to obtain the remote WWNN information.

-type fcp | escon
(Optional) The type of the connection over which the path was created.

fcp Fibre-channel protocol

escon Enterprise Systems Connection (IBM Z).

-quiet
(Optional) Turns off the Remote Mirror and Copy path removal confirmation prompt for this command.

-force
(Optional) Specifies that you want to remove Remote Mirror and Copy paths even if Remote Mirror and Copy volume pairs exist. Otherwise, specified paths that are associated with existing Remote Mirror and Copy volume pairs are not be removed.

source_LSS_ID:target_LSS_ID ... | **-**
(Required) Specifies the Remote Mirror and Copy path relationships for the source and target LSSs that are to be removed. The LSS pair ID consists of two LSS IDs, one designated as the source LSS and the other as the target LSS for a Remote Mirror and Copy path relationship. The two LSS IDs must be separated with a colon and no spaces. The first LSS ID is the source LSS. The second LSS ID is the target LSS.

This parameter accepts fully qualified LSS IDs, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified.

The fully qualified LSS ID format is *storage_image_ID/XX*, where *XX* is a hexadecimal number in the range 00 - FE.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified pair: IBM.2107-75FA120/00:IBM.2107-75FA150/00

Example of a shortened version: 00:00

Example of multiple pairs: 00:00 01:01 02:02

Example

Invoking the rmpprcpath command

```
dscli> rmpprcpath -dev IBM.2107-75FA120
-remotedev IBM.2107-75FA150 -remotewwnn 12341234000A000F 01:01
```

The resulting output

```
Are you sure want to remove the Remote Mirror and Copy path 01:01?
 [y/n]: Y
Remote Mirror and Copy path 01:01 successfully removed.
```

Remote Mirror and Copy commands

Remote Mirror and Copy (formerly PPRC) commands are used to create, manage, view, and delete Remote Mirror and Copy pairs.

The following Remote Mirror and Copy pair commands are available:

failbackpprc
Copies the required data from the source volume to the target volume to resume mirroring. You can use this command after a **failoverpprc** command has been issued to restart mirroring from site A (local site) to site B (remote site).

failoverpprc

Generates Global Mirror and Metro Mirror disaster recovery processes with the following results:

- In a Global Mirror failover recovery process, the **failoverpprc** command initiates failover processing of B volumes to A volumes.
- In a Global Mirror failback recovery process (production is returned to the local site), the **failoverpprc** command initiates failover processing from A volumes to B volumes.
- In a Metro Mirror disaster recovery process, failover processing to the Global Copy secondary volume causes the secondary volumes to become primary volumes and immediately suspends these volumes. The **failoverpprc** command changes a secondary device into a primary suspended device while leaving the primary device in its current state.

lspprc Generates a report that displays a list of remote mirror and copy volume relationships for a storage image and the status information for each remote mirror and copy volume relationship in the list.

mkpprc Establishes a remote mirror and copy relationship for a volume pair.

chpprc Modifies the characteristics of an existing Remote Mirror and Copy relationship.

freezepprc
Creates a new remote mirror and copy consistency group. It places the source logical subsystem (LSS) in the long busy state so that no I/O can be directed to it. It also removes remote mirror and copy paths between the source LSS and target LSS and sets the queue-full condition for the primary volume.

pausepprc
Pauses an existing remote mirror and copy volume pair relationship or pauses a single volume ID.

resumepprc
Resumes a remote mirror and copy relationship for a volume pair.

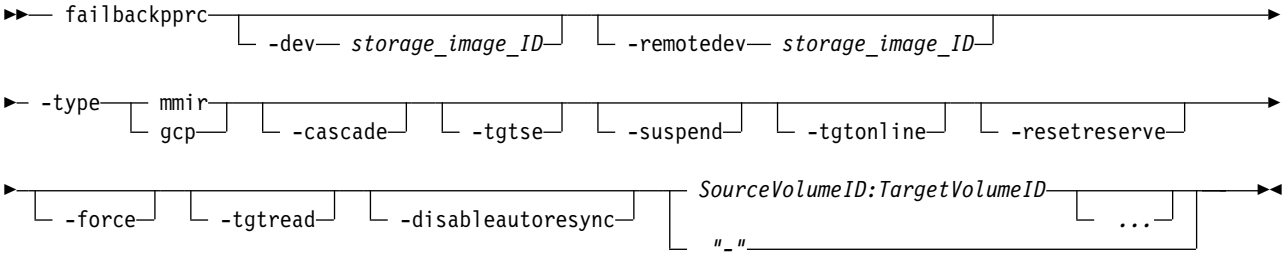
rmpprc Removes one or more specified remote mirror and copy volume pair relationships, or it removes a single volume ID (which might be useful when a disaster occurs and you want to specify only the available volume and not both the primary and secondary volumes).

unfreezepprc
Resumes I/O activity on a storage system where the **freezepprc** command has been issued. The **unfreezepprc** command resets the queue full condition for the primary volume.

failbackpprc

The **failbackpprc** command copies the required data from the source volume to the target volume to resume mirroring.

This command is used in the disaster recovery processes that are associated with sites using Metro Mirror, Global Mirror, or Metro/Global Mirror processing.



Parameters

Notes:

1. You can enter the **failbackpprc** command against any remote mirror and copy volume that is in a primary suspended state. The failback processing copies the required data from the source volume to the target volume to resume mirroring.
2. A Metro Mirror (synchronous) pair must be suspended before it can be reestablished as a Global Copy (extended distance) pair and conversely the opposite way.
3. When you specify subsystem IDs (SSIDs), the source and target volumes are restricted to 1 LSS for the source and 1 LSS for the target.

-dev *storage_image_ID*
(Optional) Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you don't specify a fully qualified source volume ID (which includes the storage image ID, for the source volume IDs that are defined by the **Source_Volume_ID:Target_Volume_ID** parameter), don't set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Note: The use of the **failbackpprc** command requires a role reversal for this parameter. The value for this parameter must be the **original primary** site that has been repaired and is ready to once again become your primary production site. For example:

- Original primary site (Site A) has a value of IBM.2107-75FA120 with volumes 0100, 0101, 0102, 0103.

- Original secondary site (Site B) has a value of IBM.2107-75FA150 with volumes 0200, 0201, 0202, 0203.
- The following **failbackpprc** command is correct:

  dscli> failbackpprc -dev IBM.2107-75FA120
  -remotedev IBM.2107-75FA150 0100:0200 0101:0201 0102:0202 0103:0203

-remotedev *storage_image_ID*
(Optional) Specifies the target storage image ID, which includes manufacturer, type, and serial number. The **-remotedev** parameter identifies the remote storage system that contains the target volume IDs that are defined by the **SourceVolumeID:TargetVolumeID** parameter. The **-remotedev** parameter is required if you don't specify a fully qualified target volume ID or if you use the **-dev** parameter.

Note: The use of the **failbackpprc** command requires a role reversal for this parameter. The value for this parameter must be the **original secondary** site. For example:

- Original primary site (Site A) has a value of IBM.2107-75FA120 with volumes 0100, 0101, 0102, 0103.
- Original secondary site (Site B) has a value of IBM.2107-75FA150 with volumes 0200, 0201, 0202, 0203.
- The following **failbackpprc** command is correct:

  dscli> failbackpprc -dev IBM.2107-75FA120
  -remotedev IBM.2107-75FA150 0100:0200 0101:0201 0102:0202 0103:0203

-type mmir | gcp
(Required) Modify one or more existing remote mirror and copy volume relationships as either Metro Mirror (Synchronous) or Global Copy (Extended Distance) relationships.

mmir Metro Mirror maintains the remote mirror and copy relationship in a consistent manner by returning the I/O write completion status to the application when the updates are committed to the target. This process becomes slower as the physical distance between source and target increases.

gcp Global Copy maintains the remote mirror and copy relationship in a nonsynchronous manner. I/O write completion status is returned to the application when the updates are committed to the source. Updates to the target volume are completed later. The original order of updates isn't strictly maintained.

-cascade
(Optional) Specifies that the remote mirror and copy target volume can also be a remote mirror and copy source volume of a different remote mirror and copy volume relationship.

-tgtse
(Optional) Specifies that the PPRC secondary volume is a space efficient volume.

-suspend
(Optional) Specifies that the remote mirror and copy relationship is to be suspended when the task completes.

Notes:

1. This parameter isn't valid for a Global Copy (Extended Distance) remote mirror and copy volume relationship.
2. This parameter isn't valid for a Metro Mirror (Synchronous) remote mirror and copy volume relationship that is established with the No Copy option activated.

-tgtonline
(Optional) Specifies that a remote mirror and copy volume relationship is to be established, including when the target volume is online to host systems.

This parameter applies to IBM Z volumes. Itdoesn't apply to open systems volumes.

-resetreserve
(Optional) Specifies that a remote mirror and copy relationship is to be established when the volume on the secondary logical subsystem is reserved by another host. If this parameter isn't specified and the volume on the secondary logical subsystem is reserved, the command fails.

Note: This parameter applies only to fixed block volumes.

-force
(Optional) Specifies whether validation of the volumes involved in the establish request occurs or is bypassed. This parameter allows you to create a Metro Mirror pair between two volumes who had no previous relationship and *ONLY* copy changed tracks.

Notes:

1. This parameter can only be used as part of a Metro/Global Mirror (3-site) disaster recovery process.
2. Only use this parameter if you're fully aware of the affect this parameter has on your transactions. A couple of scenarios are provided in this guide that describe a set of circumstances that allows you to safely use this parameter. If your circumstances don't match the scenarios, you're cautioned not to use this parameter unless advised to do so by your technical support representative.

-tgtread
(Optional) Specifies that host servers be allowed to read from the remote mirror and copy target volume. For a host server to read the volume, the remote mirror and copy pair must be in a full-duplex state.

Note: This parameter applies only to Open System volumes.

-disableautoresync
(Optional) Allows you to disable the mechanism that automatically resumes a suspended Global Copy relationship. The default isn't disabled.

SourceVolumeID:TargetVolumeID ... | -
(Required) Specifies the remote mirror and copy volume pair IDs for the source and target volume pairs that are to undergo failback processing. The original values (before the disaster) return with the source volume IDs equal to the A volumes and the target volume IDs equal to the B volumes.

This parameter accepts fully qualified volume IDs, which includes storage image IDs or a shortened version without storage image IDs if the **-dev** parameter is specified. You must separate one or more remote mirror and copy pair IDs with spaces.

A remote mirror and copy volume pair ID consists of two volume IDs. For a remote mirror and copy relationship, one volume is identified as the source and the other as the target volume. You must separate the two volume IDs of a remote mirror and copy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

X Specifies the address group, 0 - F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

The ellipsis (...) indicates that, optionally, you can specify one or more values. If you use the dash (-), the specified value is read from standard input. You can't use the dash (-) while you're in the DS CLI interactive command mode.

Example of a fully qualified pair: IBM.2107-75FA120/0100:IBM.2107-75FA150/0100

Example of multiple pairs: 0101:0101 0102:0102 0103:0103

Example

Invoking the failbackpprc command

```
dscli> failbackpprc -dev IBM.2107-75FA120
-remotedev IBM.2107-75FA150 0100:0100 0101:0101 0102:0102 0103:0103
```

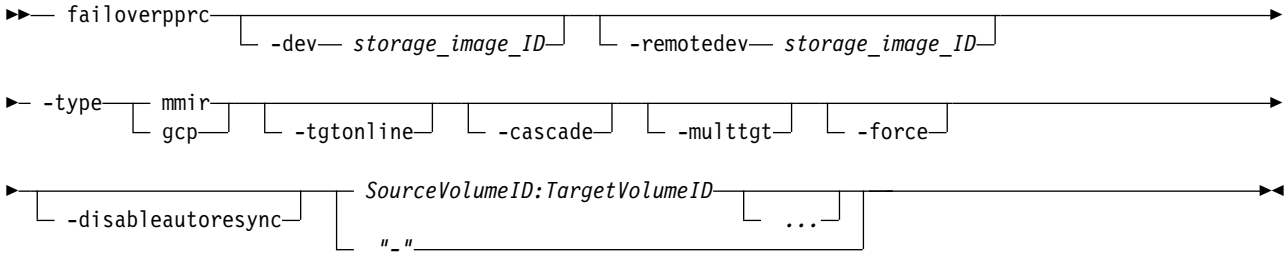
The resulting output

```
Remote Mirror and Copy pair IBM.2107-75FA120/0100:IBM.2107-75FA150/0100
 successfully failed back.
Remote Mirror and Copy pair IBM.2107-75FA120/0101:IBM.2107-75FA150/0101
successfully failed back.
Remote Mirror and Copy pair IBM.2107-75FA120/0102:IBM.2107-75FA150/0102
 successfully failed back.
Remote Mirror and Copy pair IBM.2107-75FA120/0103:IBM.2107-75FA150/0103
 successfully failed back.
```

failoverpprc

The **failoverpprc** command is used only with disaster recovery processing.

This command is used in the disaster recovery processes associated with sites that use Metro Mirror, Global Mirror, or Metro/Global Mirror processing. The **failoverpprc** command succeeds even if the paths are down and the volume at the production site is unavailable or nonexistent.



Parameters

The **failoverpprc** command is used in the Global Mirror and Metro Mirror disaster recovery processes with the following results:

- In a Global Mirror failover recovery process, the **failoverpprc** command initiates failover processing of B volumes to A volumes. This process causes the B volumes to become the primary volumes and the A volumes to become the secondary volumes. The effect is that the Global Copy state of the B volumes changes from secondary to primary and suspended.
- In a Global Mirror failback recovery process (production is returned to the local site), the **failoverpprc** command initiates failover processing from A volumes to B volumes. This process causes the A volumes to become the primary volumes and the B volumes to become the secondary volumes.
- In a Metro Mirror disaster recovery process, failover processing to the Global Copy secondary volume causes the secondary volumes to become primary volumes and immediately suspends these volumes. When you run a Global Copy failover, the B volumes are the primary volumes and the A volumes are the secondary volumes. This action changes only the Global Copy state of the secondary volumes from Target Copy Pending to Suspended. The **failoverpprc** command changes a secondary device into a

primary suspended device while leaving the primary device in its current state. This command succeeds even if the paths are down and the volume at the production site is unavailable or nonexistent.

Note: When you specify the subsystem identifier (SSID), the source and target volumes are restricted to one LSS for the source and one LSS for the target.

-dev *storage_image_ID*
(Optional) Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified source volume ID, do not set the *devid* variable in your profile or through the **setenv** command. The storage image ID is also required if the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Note: The use of the **failoverpprc** command requires a role reversal for this parameter. The value for this parameter must be the **original secondary** site. For example:

- Original primary site (Site A) has a value of IBM.2107-75FA120 with volumes 0100, 0101, 0102, 0103.
- Original secondary site (Site B) has a value of IBM.2107-75FA150 with volumes 0200, 0201, 0202, 0203.
- The following **failoverpprc** command is correct:

```
dscli> failoverpprc -dev IBM.2107-75FA150
-remotedev IBM.2107-75FA120 0200:0100 0201:0101
0202:0102 0203:0103
```

-remotedev *storage_image_ID*
(Optional) Specifies the target storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified target volume ID or if you use the **-dev** parameter.

Note: The use of the **failoverpprc** command requires a role reversal for this parameter. The value for this parameter must be the **original primary** site. For example:

- Original primary site (Site A) has a value of IBM.2107-75FA120 with volumes 0100, 0101, 0102, 0103.
- Original secondary site (Site B) has a value of IBM.2107-75FA150 with volumes 0200, 0201, 0202,0203.
- The following **failoverpprc** command is correct:

```
dscli> failoverpprc -dev IBM.2107-75FA150
-remotedev IBM.2107-75FA120 0200:0100 0201:
0101 0202:0102
```

-type mmir | gcp
(Required) Modifies one or more existing remote mirror and copy volume relationships as either Metro Mirror or Global Copy relationships.

mmir Metro Mirror maintains the remote mirror and copy relationship in a consistent synchronous manner when the updates are committed to the target. This process becomes slower as the physical distance between source and target increases.

gcp Global Copy maintains the remote mirror and copy relationship in a nonsynchronous manner when the updates are committed to the source. Updates to the target volume are done later. The original order of updates is not strictly maintained.

-tgtonline
(Optional) Establishes a remote mirror and copy volume relationship, including when the target volume is online to host systems.

This parameter applies to IBM Z volumes. It does not apply to open systems volumes.

-cascade
(Optional) Specifies that the PPRC target volume can also be a PPRC source volume of a different PPRC volume relationship.

-multtgt
(Optional) Specifies to convert to a multi-target relationship when failing over the A - B relationship in a cascading A – B – C setup configuration.

-force
(Optional) Specifies whether validation of the volumes that are involved in the establish request occurs or is bypassed. With this parameter, you can create a PPCR pair between two volumes that have no previous relationship and only changed tracks are copied.

Notes:

1. When you specify the **-force** parameter, PPRC paths must be established between the source volume and the target volume.
2. This parameter can be used only as part of a Metro/Global Mirror (three-site) disaster recovery process.
3. Use the **-force** parameter if you are fully aware of the affect that this parameter has on your transactions. A couple of tasks are provided as examples to help you safely use this parameter. If your circumstances do not match the scenarios, you are cautioned not to use this parameter unless advised to do so by your technical support representative.

-disableautoresync
(Optional) Disables the mechanism that automatically resumes a suspended Global Copy relationship. The default is not disabled. The **-disableautoresync** parameter is available only in Version 5 Release 3 or later.

SourceVolumeID:TargetVolumeID ... | –
(Required) Specifies the remote mirror and copy volume pair IDs of the source and target volumes that must have their relationships changed so that the target volumes (B volumes) become the source volumes and the original source volumes (A volumes) become the target volumes. This results in the following conditions:

- The source volumes (B volumes) show as a suspended host.
- The target volumes (A volumes) show as a suspended target and they are accessible for mounting.

This parameter also accepts fully qualified volume IDs, which include storage image IDs or a shortened version without storage image IDs, if the **-dev** parameter is specified. You must separate multiple remote mirror and copy pair IDs with spaces.

A remote mirror and copy pair ID consists of two volume IDs: one designated as the source and the other as the target volume for a remote mirror and copy relationship. You must separate the two volume IDs of a remote mirror and copy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

X Specifies the address group, 0 - F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified pair: IBM.2107-75FA150/0100:IBM.2107-75FA120/0100

Example of a shortened version: 0100:0100

Example of multiple pairs: 0101:0101 0102:0102 0103:0103

Example

Invoking the failoverpprc command

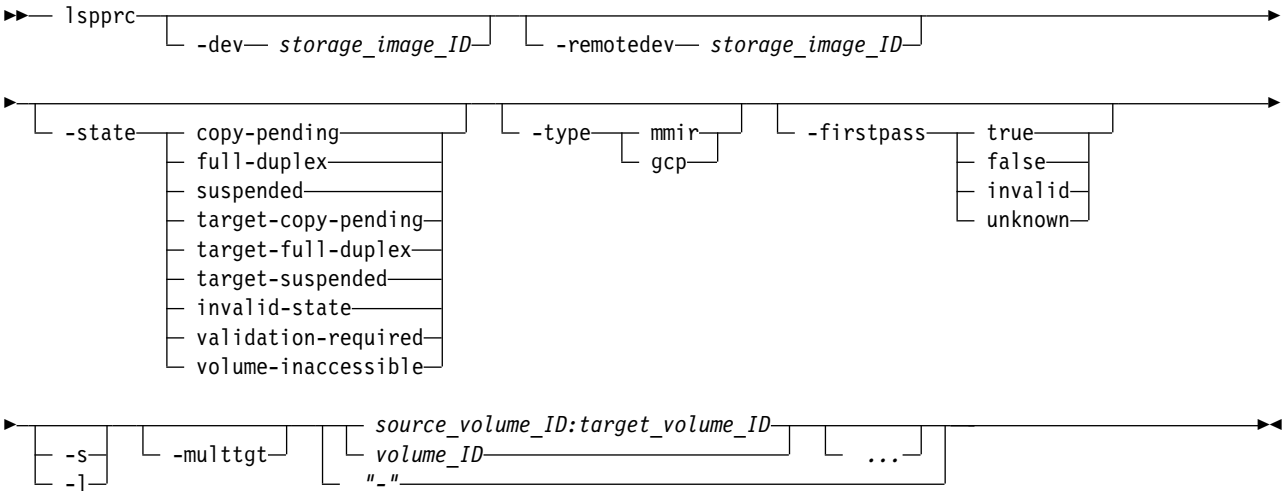
```
dscli> failoverpprc -dev IBM.2107-75FA150
-remotedev IBM.2107-75FA120 0200:0100 0201:0101 0202:0102
```

The resulting output

A confirmation message is presented for each remote mirror and copy pair that is successfully suspended.

lspprc

The **lspprc** command displays a list of remote mirror and copy (formerly PPRC) volume relationships for a storage image. The command also displays status information for each remote mirror and copy volume relationship in the list.



Parameters

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the source and target volumes, do not set the *devid* variable in your profile or through the **setenv** command. The storage image ID is also required if the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-remotedev *storage_image_ID*
(Optional; however, required as noted). Specifies the target volume storage image ID, which consists of manufacturer, machine type, and serial number.

Note: The **-remotedev** parameter is required when volume pairs are specified and the **-dev** parameter is specified, as well.

-state copy-pending | full-duplex | suspended | target-copy-pending | target-full-duplex | target-suspended | invalid-state | validation-required | volume-inaccessible
(Optional) Specifies the state of the remote mirror and copy relationships that you want to view.

copy-pending
Specifies that you want to view remote mirror and copy relationships that have copy processing that is pending. A Global Copy relationship is always copy-pending.

full-duplex
Specifies that you want to view remote mirror and copy relationships that are full-duplex.

suspended
Specifies that you want to view remote mirror and copy relationships that are suspended. The Reason attribute might indicate why the relationship is suspended.

target-copy-pending
Specifies that you want to view remote mirror and copy relationships where the target volume has copy processing that is pending. In this state, the source volume is unknown or cannot be queried.

target-full-duplex
Specifies that you want to view remote mirror and copy relationships where the target volume is full-duplex. In this state, the source volume is unknown or cannot be queried.

target-suspended
Specifies that you want to view remote mirror and copy relationships where the target volume is suspended. In this state, the source volume is unknown or cannot be queried. The Reason attribute might indicate why the relationship is suspended.

invalid-state
Specifies that you want to view remote mirror and copy relationships with an Invalid State. The Invalid State means that a general internal error occurred when the query was processed. The report that is generated with this query displays only the source and target volume IDs of a remote mirror and copy volume relationship. The report also displays the state designation of Invalid State. All the other information columns are displayed with a " - " value.

validation-required
Specifies that further validation is required. Running the **lspprc** command again might display a different state output. However, if you request a query with just the *validation-required* designation, the report displays only the source and target volume IDs of a remote mirror and copy volume relationship that have a designation of validation-required. All the other information columns are displayed with a " - " value.

volume-inaccessible
Specifies that you want to view remote mirror and copy relationships where the volume cannot be viewed. In this case, it means that the volume is fenced. The report that is generated with this query displays only the source and target volume IDs of a remote mirror and copy volume relationship and the state designation of *volume-inaccessible*. All the other information columns are displayed with a " - " value.

-type mmir | gcp
(Optional) Specifies the type of the remote mirror and copy relationships that you want to view.

mmir Specifies that you want to view Metro Mirror relationships.

gcp Specifies that you want to view Global Copy relationships.

-firstpass true | false | invalid | unknown
(Optional) Specifies the first pass status of the Global Copy relationships that you want to view.

true The first pass of the Global Copy is complete.

false The first pass of the Global Copy is not yet complete.

invalid
The remote mirror and copy relationship is not a Global Copy relationship, or the query was sent to the secondary volume of the Global Copy relationship.

unknown
The first pass status of the Global Copy is unknown.

-s
(Optional) Displays the remote mirror and copy volume pair IDs. You cannot use the **-s** and the **-l** parameters together.

-l
(Optional) Displays the default output plus extra attributes that are identified as long output. You cannot use the **-s** and the **-l** parameters together.

-multtgt
(Optional) Displays internal remote mirror and copy relationships that are created automatically by the DS8000 as a part of a multi-target configuration. The remote mirror and copy relationships are suspended, unless they are needed in a recovery situation to resynchronize the two volumes, without requiring a full-volume copy.

SourceVolumeID:TargetVolumeID | Volume_ID ... | -
(Required) Displays the remote mirror and copy relationships for the source and target volume pairs, or for volumes with the specified IDs.

This parameter accepts fully qualified volume IDs, which include storage image IDs or a shortened version without storage image IDs, if the **-dev** parameter is specified. You must separate multiple remote mirror and copy pair IDs with spaces.

A remote mirror and copy pair ID consists of two volume IDs: one designated as the source and the other as the target volume for a remote mirror and copy relationship. You must separate the two volume IDs of a remote mirror and copy pair ID with a colon and no spaces. The first volume ID is the source volume. The second volume ID is the target volume.

You can enter remote mirror and copy pair IDs, a range of remote mirror and copy pair IDs, single volume IDs, or a range of volume IDs. You cannot enter a combination of remote mirror and copy pair IDs and volume IDs.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a shortened version pair: 0100:0100

Example of multiple pair IDs: 0100:0100 0200:0200 0300:0300

Example of a range of pair IDs: 0100-010F:0100-010F

Example of a source or target volume ID: 0100

Example of a range of source or target volume IDs: 0100-010F

Note: A query of target volume IDs is directed to the storage image that is identified by the **-dev** parameter or embedded in the fully qualified single volume IDs.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lspprc** command with the **-l** parameter.

Invoking the lspprc command

```
dscli> lspprc -dev IBM.2107-75FA120
-remotedev IBM.2107-75FA150 -l 0100:0100 0101:0101
```

The resulting output

ID	State	Reason	Type	Out Of Sync Tracks	Tgt Read	Src Cascade
IBM.2107-75FA120/0100: IBM.2107-75FA150/0100	Full-Duplex	-	Metro Mirror	0	Enabled	Disabled
IBM.2107-75FA120/0101: IBM.2107-75FA150/0101	Full-Duplex	-	Metro Mirror	0	Enabled	Disabled

Tgt Cascade	Date Suspended	Source LSS	Timeout (secs)	Critical Mode	First Pass Status	Incremental Resync	Tgt Write
Enabled	-	01	120	Disabled	False	Enabled	Enabled
Enabled	-	02	120	Disabled	False	Enabled	Enabled

GMIR CG	PPRC CG	isTgtSE	DisableAutoResync
Disabled	Disabled	Unknown	-
Disabled	Disabled	Unknown	-

Report field descriptions

ID Indicates the source and target volume IDs of a remote mirror and copy volume relationship. Use the **-fullid** parameter to display fully qualified IDs, which include the storage image ID.

State Displays the current remote mirror and copy volume relationship state. One of the following values is displayed:

Copy Pending
Indicates that the relationship is copy pending. A Global Copy (extended distance) relationship is always copy pending.

Full Duplex
Indicates that the relationship is full duplex.

Suspended
Indicates that the relationship is suspended. The Reason attribute might indicate why the relationship is suspended.

Target Copy Pending
Indicates that the source volume is unknown or cannot be queried and the target state is copy pending.

Target Full Duplex
Indicates that the source volume is unknown or cannot be queried and the target state is full-duplex.

Target Suspended
Indicates that the source volume is unknown or cannot be queried and the target state is suspended.

Invalid State
Indicates that a general internal error occurred when the query was processed.

Note: The report that is generated with the *Invalid State* designation displays only the source and target volume IDs of a remote mirror and copy volume relationship and the state designation of Invalid State. All the other information columns are displayed with a " - " value.

Validation Required
Indicates that the status of the volume cannot be currently determined. Further validation is required. Sometimes, running the same **lspprc** command again generates different state output.

Note: The report that is generated with the *validation-required*designation displays only the source and target volume IDs of a remote mirror and copy volume relationship and the state designation of validation-required. All the other information columns are displayed with a " - " value.

Volume Inaccessible
Indicates that the volume might not be queried. This information means that the volume is fenced.

Note: The report that is generated with the *Volume Inaccessible* designation displays onlythe source and target volume IDs of a remote mirror and copy volume relationship and the state designation of *Volume Inaccessible*. All the other information columns are displayed with a " - " value.

Reason
Indicates why the remote mirror and copy volume relationship is suspended or copy-pending. The following values can be displayed:

-
Indicates that the volume is suspended but the reason for the suspension is not defined within the system.

Not in PPRC Relationship
Indicates that the designated volume is not part of a remote mirror and copy pair.

Host Source
Indicates that the remote mirror and copy processing on the volume was suspended by the primary host. The host command might specify an immediate suspension or that the volume is to be suspended when it entered a full duplex state.

Host Target
Indicates that remote mirror and copy processing was suspended on the secondary volume. Updates to primary volumes and out-of-sync tracks are still being processed.

Update Target
Indicates that remote mirror and copy processing was suspended on a secondary volume by the primary control unit update secondary device status command.

Internal Conditions Both
Indicates that remote mirror and copy processing was suspended on a volume by either the primary control unit or the secondary control unit because of internal conditions.

Simplex Target
Indicates that remote mirror and copy processing was suspended on a volume when the secondary control unit sent a state change to the primary control unit indicating a transition to a simplex state.

Internal Conditions Target
Indicates that remote mirror and copy processing was suspended on a secondary volume when the primary control unit was notified that the secondary volume became suspended due to internal conditions.

Power Indicates that remote mirror and copy processing was suspended on a volume when the primary or secondary control unit was shut down or restarted.

Notes:

1. The paths to the secondary controller might not be available if the power to the primary controller was shut down. If only the secondary control unit was shut down, it might be possible for the paths to be restored depending on the path status. Use the following process to determine whether your remote mirror and copy processing can be restored on the affected volumes:
   a. Enter the **lspprc** command and use the generated report to determine the path status.
   b. Enter the **mkpprc** command if the paths are still in tact. This process resynchronizes the volume pairs.
   c. Continue with your processing.
2. If the previous process cannot be completed, you must remove the pair relationships on the affected volumes and start your remote mirror and copy processing from the beginning on the affected volumes.

Freeze Indicates that remote mirror and copy processing was suspended on a volume pair because the host issued a Freeze PPRC Group order.

Volume Not Configured
Indicates that remote mirror and copy processing was suspended on a volume because the volume is not part of a copy pair relationship.

Remote Copy Pending
Indicates that the volume transitioned to the pending state as a result of establishing a FlashCopy relationship. Before the attempt was made to establish the FlashCopy relationship, the volume was in Full Duplex mode.

Release Space Failure
Indicates that the pair is suspended due to the failure to release unused space-efficient storage on the Remote Mirror and Copy secondary volume.

With Secondary Consistency
Indicates that the Remote Mirror and Copy secondary volumes form a consistent data set.

Type Indicates that the remote copy and mirror volume relationship is a Metro Mirror (synchronous) relationship, a Global Copy (extended distance) relationship, or the relationship type is unknown.

Out Of Sync Tracks
Indicates the number of tracks that are not synchronized for this FlashCopy relationship. The maximum value depends on the source volume size.

Notes:

1. If you enter the **lspprc** command to view the out-of-sync value for a volume pair (for example, 0000:0001) on a 2105, there is no observable decrease in the value from when you issue the query to the end of the process.

2. If you enter the **lspprc** command to view the out-of-sync value for a single volume (for example, 0000) on a 2105, there is an observable decrease in the value but only at 10-second intervals. If you enter the **lspprc** command and reissue it again before 10 seconds expires, there is no observable change in the value.
3. If you enter the **lspprc** command to view the out-of-sync value for a volume pair or a single volume, there is an observable decrease in the value but only at 10-second intervals.

Tgt Read
Indicates that Read I/O operations to the target volume are allowed.

Src Cascade
Indicates that the source volume of this relationship is enabled to also be a target volume of a different relationship.

Tgt Cascade
Indicates that the target volume of this relationship is enabled so that it is also a source volume for a different relationship.

Date Suspended
Indicates the date when this relationship was last suspended. The value can be displayed as a " - ".

SourceLSS
Indicates the consistency group LSS ID that is associated with the source volume of this PPRC volume relationship.

Timeout (secs)
Indicates the consistency group Long Busy Timeout setting for the LSS ID that is associated with the source volume of this PPRC volume relationship. This value can be modified by entering the **chlss** (FB) or the **chlcu** (CKD) command.

Critical Mode
Indicates whether the remote copy and mirror primary volume represents a critical volume.

First Pass Status
Indicates whether the first pass of Global Copy is complete. When you query the primary volume of a Global Copy pair, *True* is displayed when the first pass is complete, *False* is displayed when the first pass is not yet complete, and *Invalid*is displayed if the pair is not Global Copy or if the secondary was queried.

Incremental Resync
Indicates whether incremental resynchronization is running.

Tgt Write
Indicates whether input is allowed to the remote mirror and copy secondary volume.

GMIR CG
The value for this field is fixed to display N/A (not available). Use the **lssession** command to monitor the status of Global Mirror relationships.

PPRC CG
Indicates whether the remote mirror and copy consistency group is enabled, disabled, or not available.

Notes:

1. This value is displayed when you designate the use of the **-l** parameter, and when the primary volume is being queried. If the secondary volume is being queried, the value that is displayed for this field is N/A (not available).
2. This value is not reported for model 2105. If a model 2105 is being queried, the value that is displayed for this field is N/A (not available).

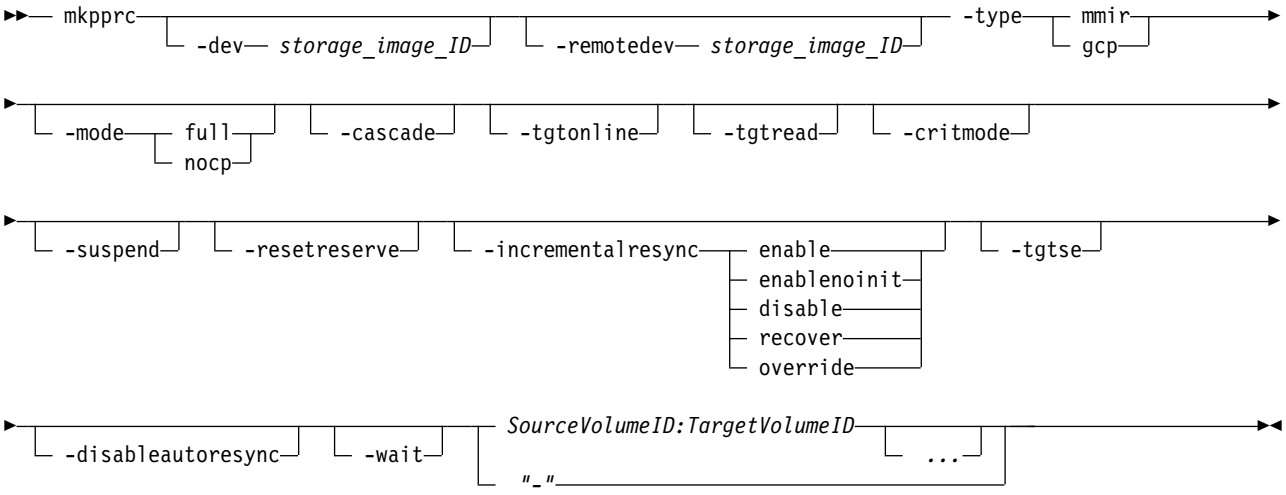
isTgtSE
Indicates whether this remote mirror and copy relationship has a space-efficient secondary.

Unknown
Indicates that the space allocation method of the target is not known.

DisableAutoResync
Indicates whether the mechanism that automatically resumes a suspended Global Copy relationship is active.

mkpprc

The **mkpprc** command establishes a remote mirror and copy (formerly PPRC) relationship for a volume pair.



Parameters

Notes:

1. When you specify subsystem IDs, the source and target volumes are restricted to one LSS for the source and one LSS for the target.
2. If you are using the Cisco MDS 9216 Multilayer Fabric Switch, you must not enable the write acceleration feature. The **mkpprc** command might fail if the write acceleration feature is enabled.

-dev *storage_image_ID*
(Optional) Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified source volume ID, do not set the *devid* variable in your profile or through the **setenv** command. The storage image ID is also required if the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-remotedev *storage_image_ID*
(Optional) Specifies the target storage image ID, which includes manufacturer, machine type, and serial number. This parameter is required if you do not specify a fully qualified target volume ID, or if the **-dev** parameter is selected.

-type mmir | gcp
(Required) Establishes one or more remote mirror and copy volume relationships as either Metro Mirror (synchronous) or Global Copy (extended distance) relationships.

mmir Metro Mirror maintains the remote mirror and copy relationship in a consistent

(synchronous) manner by returning I/O write completion status to the application when the updates are committed to the target. This process becomes slower as the physical distance between source and target increases.

gcp Global Copy maintains the remote mirror and copy relationship in a nonsynchronous manner. I/O write completion status is returned to the application when the updates are committed to the source. Updates to the target volume are performed at a later time. The original order of updates is not strictly maintained.

-mode full | nocp
(Optional) Specifies the following initial data copy mode for the remote mirror and copy volume relationships:

full Full mode copies the entire source volume to the target volume. This value is the default when you do not specify the no copy mode.

nocp No Copy mode does not copy data from source to target volumes. This option presumes that the volumes are already synchronized.

-cascade
(Optional) Enables a remote mirror and copy target volume to be a remote mirror and copy source volume for a different remote mirror and copy volume relationship. The default value for this parameter is disabled.

-tgtonline
(Optional) Establishes a remote mirror and copy volume relationship, including when the target volume is online to host systems. This parameter applies to IBM Z volumes and does not apply to open systems volumes. The default value for this parameter is disabled.

-tgtread
(Optional) Allows host servers to read from the remote mirror and copy target volume. For a host server to read the volume, the remote mirror and copy pair must be in a full-duplex state. This parameter applies to open systems volumes and does not apply to z System volumes. The default value for this parameter is disabled.

-critmode
(Optional) Protects the source volume from receiving new data. If the last path fails between the pairs and results in the inability to send information to the target, the source is protected. Current updates and subsequent attempts to update the source fail with a unit check on IBM Z. The default value for this parameter is disabled.

Note: This parameter applies only to IBM Z volumes.

Critical mode operates in one of three ways that depend on the setting of the LCU critical mode and the setting of the **-critmode** parameter in this command. The Critical volume mode definition overview table describes how the critical volume mode works.

Table 11. Critical volume mode definition overview

Critical Mode	LCU, Critical Heavy	Mkpprc critmode	Description
Normal	Disabled or Enabled	Disabled	• Suspends the primary volume. • Allows write operations to the primary volume.
Critical Volume	Disabled	Enabled	• Suspends the primary volume when the last path to the secondary volume fails. • Inhibits write operations to the primary volume.

Table 11. Critical volume mode definition overview (continued)

Critical Mode	LCU, Critical Heavy	Mkpprc critmode	Description
Critical Heavy	Enabled	Enabled	• Suspends the primary volume when the secondary volume cannot be updated for any reason. • Inhibits write operations to the primary volume.

Notes:

1. Use the **-critmode** parameter only for log devices, not for devices that the system requires. In extreme cases, the host system might require you to load the initial program to recover a device that is write inhibited. Whenever possible, use the **freezepprc** command as an alternative to using the **-critmode** parameter.
2. The **-critmode** parameter cannot be used with Global Copy or remote copy and mirror cascading volumes.
3. To reset a volume that is write inhibited because of critical mode, you can enter the **mkpprc**, **pausepprc**, or **rmpprc** command to this volume.
4. Use automation software as part of any solution that includes critical mode. Automation software is not a part of the software that is provided with the storage system; you must supply it. However, IBM has offerings to assist with this automation. For more information, contact your IBM storage representative.

-suspend
(Optional) Suspends the remote mirror and copy relationship when the task completes. This parameter is not valid for a Global Copy (extended distance) remote mirror and copy volume relationship. This parameter is not valid for a Metro Mirror (synchronous) remote mirror and copy volume relationship that is established with the No Copy option. The default value for this parameter is disabled.

-resetreserve
(Optional - open system volumes only) Allows the establishment of a remote mirror and copy relationship when the volume on the secondary logical subsystem is reserved by another host. You can use this parameter only with open system volumes. If this option is not specified and the volume on the secondary logical subsystem is reserved, the command fails.

-incrementalresync enable | enablenoinit | disable | recover | override
(Required for three-site Metro/Global Mirror recovery) Specifies the resynchronization method that is used in a three-site Metro/Global Mirror disaster recovery configuration. A three-site Metro/Global Mirror configuration usually involves the following sites:

- Site A (local site), which contains the production volumes (or the A volumes).
- Site B (intermediate site), which contains the B volumes and a local, synchronous copy.
- Site C (remote site), which contains the C volumes and an asynchronous copy that is located remotely from sites A and B.

You can specify the following options when you first establish volume pairs by using the **mkpprc** command or on established volume pairs with this command.

enable
Specifies that change recording features be created on each of the primary Metro Mirror volumes to enable the microcode to monitor and track data in a Metro/Global Mirror configuration. This is data that has been written to the primary volumes but not secured on the remote site volumes.

enablenoinit
Specifies that the change recording features are not created or started on the primary Metro Mirror volumes in a Metro/Global Mirror configuration. Only use this option in specific recovery scenarios.

disable
Specifies that the current incremental resynchronization function be stopped on the primary volumes of the Metro Mirror volume pairs.

recover
Specifies that a new Global Mirror volume pair be established by using an existing primary Metro Mirror (A) volume at the local site and a secondary Global Copy (C) volume at the remote site. When this command processes, only changes to the local Metro Mirror volumes are copied to the volumes at the remote site.

Note: When you specify this option, the system verifies that the failed secondary volumes for the volumes in a Metro Mirror relationship are the primary volumes for the volumes in a Global Copy relationship and that the specified volumes have the intermediate volumes in common. That is, the primary specified volumes are the A volumes and the secondary specified volumes are the C volumes in a Metro/Global Mirror configuration.

override
Specifies that a new Global Mirror volume pair be established by using an existing primary Metro Mirror (A) volume at the local site and a secondary Global Copy (C) volume at the remote site. When this command processes, only changes to the local Metro Mirror volumes are copied to the volumes at the remote site.

Note: When you specify this option, there is no validation to ensure the secondary relationship in this configuration before the **mkpprc** command is processed. Therefore, you must ensure that the primary specified volumes are the A volumes and the secondary specified volumes are the C volumes in a Metro/Global Mirror configuration.

-tgtse
(Optional) Specifies that the secondary volume might be a space-efficient logical volume.

-disableautoresync
(Optional) Disables the mechanism that automatically resumes a suspended Global Copy relationship. The default is not disabled. The **-disableautoresync** parameter is available only in Version 5 Release 3 or later.

-wait
(Optional) Delays the command response until the remote copy and mirror volumes are in one of the final states: simplex, full duplex, suspended, secondary full duplex, secondary suspended, or secondary simplex (until the pair is not in the Copy Pending state). The default value for this parameter is disabled.

Notes:

1. This parameter cannot be used with the **-type** *gcp* or **-mode** *nocp* parameters.
2. When you use the **-wait** parameter, you must wait until the original command completely processes before you can enter a new command.
3. If you are running in single-shot mode, you can periodically enter the **lspprc** command to check the remote mirror and copy volume pair state. Then, continue with new commands until the correct state is reached.

SourceVolumeID:TargetVolumeID ... | -
(Required) Specifies the remote mirror and copy volume relationship for the source and target volume pairs with the specified IDs.

This parameter accepts fully qualified volume IDs, which include storage image IDs or a shortened version without storage image IDs, if the **-dev** parameter is specified. You must separate multiple remote mirror and copy pair IDs with spaces.

A remote mirror and copy pair ID consists of two volume IDs: one designated as the source and the other as the target volume for a remote mirror and copy relationship. You must separate the two volume IDs of a remote mirror and copy pair IDs with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a fully qualified pair: IBM.2107-75FA120/0100:IBM.2107-75FA150/0100

Example of a shortened version: 0100:0100

Example of multiple pairs: 0101:0101 0102:0102 0103:0103

Extra processing tips

The following examples represent some CLI volume format options that you might want to incorporate in your processing:

Processing multiple volumes

The following two examples are ways that this environment might be processed, and both are correct. The first method might be fine if you were managing a few volumes, while the second processes hundreds or thousands of volumes more efficiently.

```
mkpprc -dev IBM.2107-13AB79A -remotedev IBM.2107-13AB76A
–type mmir -mode full
-tgtread 1000:1205 1001:1206 1002:1207 1003:1208  1004:1209 ... and so on
mkpprc -dev IBM.2107-13AB79A -remotedev IBM.2107-13AB76A
–type mmir -mode full -tgtread 1000-1004:1205-1209
```

Using the grouping method in your command

You can also group the volumes. However, the order of the volumes is critical when you group them, and they must be contiguous. The following example shows how to code for grouping:

```
mkpprc -dev IBM.2107-13AB79A -remotedev IBM.2107-13AB76A
–type mmir -mode full
-tgtread 1000-1004:1205-1209 1100-1104:1300-1304
```

Example

Invoking the mkpprc command

```
dscli> mkpprc -dev IBM.2107-75FA120
-remotedev IBM.2107-75FA150 0100:0100 -type mmir 2100-2107:2100-2107
```

The resulting output (an example only)

```
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
 2100:2100 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
 2101:2101 successfully created.
```

```
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
 2102:2102 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
 2103:2103 successfully created.
```

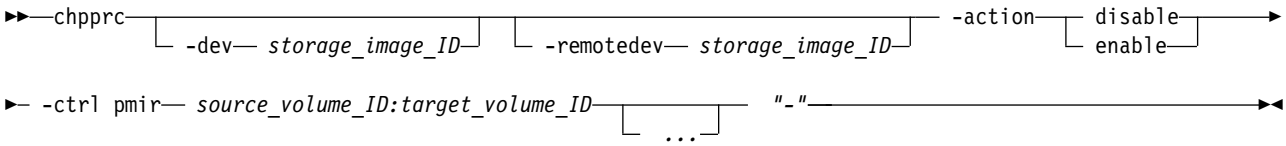
```
dscli> mkpprc -dev IBM.2107--75FA120
-remotedev IBM.2107--75FA150 -type gcp -incrementalresync enablenoinit
-mode nocp 2100-2107:2100-2107
```

The resulting output (an example only)

```
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
2100:2100 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
2101:2101 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
2102:2102 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
2103:2103 successfully created.
```

chpprc

The **chpprc** command modifies the characteristics of an existing Remote Mirror and Copy relationship.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID. A storage image ID consists of one or more source volumes IDs that are defined by the *source_volume_ID:target_volume_ID* variable. For example, DS8000 models use IBM.2107-75FA120.

-remotedev *storage_image_ID*
(Required) Specifies the remote storage image ID, which consists of a manufacturer, machine type, and serial number. The remote storage ID is required if you specify **-dev**. The **-remotedev** identifies the remote DS storage system that contains one or more target IDs that are defined by the *source_volume_ID:target_volume_ID* variable. For example, DS8000 models use IBM.2107-75FA120.

-action disable | enable
(Required) Specifies the action to apply for the specified control. The valid values are as follows:

disable
Disables a multi-target setup that is used for the Remote Pair Copy FlashCopy relationship (also known as a preserve mirror relationship).

enable
Enables a multi-target setup that is used for the Remote Pair Copy FlashCopy relationship (also known as a preserve mirror relationship).

-ctrl pmir*source_volume_ID:target_volume_ID ...* **|** **-**
(Required) Establishes a Remote Mirror and Copy volume relationship for the source and target volumes with one or more Remote Mirror and Copy IDs that you specified. Multiple Remote Mirror and Copy Pair IDs must be separated with a white space between each value.

This parameter accepts fully qualified volume IDs, which consist of storage image IDs or a shortened version without storage image IDs, if you specify the **-dev** parameter.

A Remote Mirror and Copy Pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a Remote Mirror and Copy relationship. The two volume IDs of a Remote Mirror and Copy Pair ID must be separated with a colon and no space. The first volume ID is the designated source volume. The second volume ID is the designated target volume.

The format of the volume ID can be represented as *XYZZ*, where:

X Specifies the address group, 0 - F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

An example of a shortened version is `0001:0004`

An example of multiple pairs is `0001:0004 0003:00FF 0008:000C`

Example

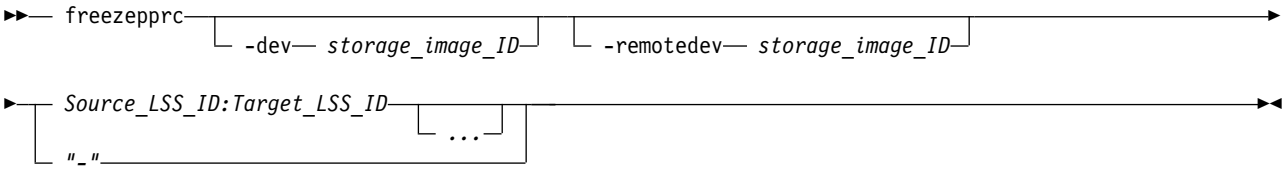
An invocation example

```
dscli> chpprc -dev IBM.2107-75FA120
```

freezepprc

The **freezepprc** command creates a new remote mirror and copy consistency group.

It places the source logical subsystem (LSS) in the *long busy* state so that no I/Os can be directed to it. It also removes remote mirror and copy paths between the source LSS and target LSS and sets the *queue full* condition for the primary volume. This causes the host to queue writes to the primary volume until the *queue full* condition is reset. During the *queue full* condition, the primary volume reports *long busy* status.



Parameters

Note: When specifying SSIDs, the command is limited to one LSS pair.

-dev *storage_image_ID*
(Optional). Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-remotedev *storage_image_ID*
(Optional). Specifies the target storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified target LSS ID or if the **-dev** parameter is used.

Source_LSS_ID:Target_LSS_ID ... | -
(Required). Specifies that a consistency group for the source and target LSSs with the IDs specified be

placed in a long busy state. Accepts fully qualified LSS IDs, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified.

A remote mirror and copy path LSS pair ID consists of two LSS IDs, one designated as the source LSS and the other as the target LSS for a remote mirror and copy path relationship. The two LSS IDs must be separated with a colon and no spaces. The first LSS ID is the source LSS. The second LSS ID is the target LSS.

The fully qualified LSS ID format is *storage_image_ID/xx*, where *xx* is a hexadecimal number in the range 00 - FE.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a pair: 00:00

Example of multiple pairs: 00:00 01:01 02:02

Example

Invoking the freezepprc command

```
dscli> freezepprc -dev IBM.2107-75FA120
-remotedev IBM.2107-75FA150 01:01
```

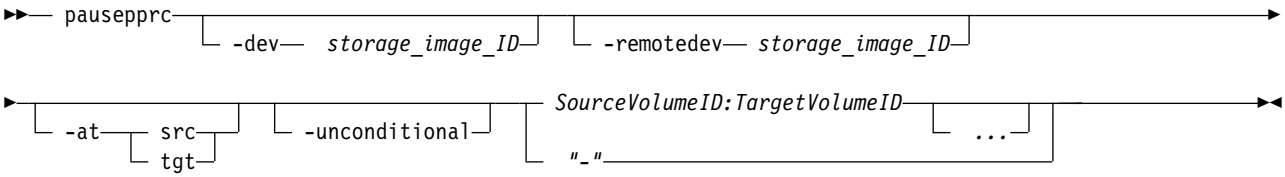
The resulting output

```
Remote Mirror and Copy consistency group 01:01 successfully created.
```

pausepprc

The **pausepprc** command pauses an existing Remote Mirror and Copy volume pair relationship or pauses a single volume ID.

To use this command with a single volume, you must specify either the **-at** *src* parameter option or the **-at** *tgt* parameter option. If neither of these options are specified in the command, single volumes are not valid.



Parameters

Note: When you specify SSIDs, the command is limited to one LSS pair.

-dev *storage_image_ID*
(Optional). Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified source volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Example of a fully qualified storage image ID: IBM.2107-75FA120.

-remotedev *storage_image_ID*
(Optional most times). Specifies the target storage image ID, which includes manufacturer, machine type, and serial number.

Note: The **-remotedev** parameter is required when volume pairs are specified and the **-dev** parameter is specified, as well.

-at src | tgt
(Optional). Specifies that you want to initiate a suspend action from either the source volume or the target volume.

src Initiates a suspend action from the source volume. After the task successfully runs, the source and target volumes are in the **suspend** state.

The **-at** *src* parameter option can also be used with single volumes. When you specify a single volume using this option, the volume is treated as a source and the target is treated as a null.

tgt Initiates a suspend action from the target volume. After the command successfully runs, the target volume is in the **suspend** state, but there is no guarantee that the source volume is suspended as well. For a suspend action that is issued to the target, the source can remain in the **active** state.

The **-at** *tgt* parameter option can also be used with single volumes. When you specify a single volume using this parameter option, the volume is treated as a target and the source is treated as a null.

Notes:

1. If you specify the **-at** *tgt* or **-at** *src* parameter and the **-unconditional** parameter, the value for the **-remotedev** parameter is ignored.
2. If you specify the **-at** *tgt* parameter and do not specify the **-unconditional** parameter, the values for the **-dev** and *SourceVolumeID* parameters are ignored.
3. If you specify the **-at** *src* parameter and do not specify the **-unconditional** parameter, the values for the **-remotedev** and *TargetVolumeID* parameters are ignored.

-unconditional
(Optional). Specifies that a source or target volume has been selected individually, and not as a pair. The **-unconditional** parameter is valid only if the **-at** parameter is selected. Do not use volume pair IDs.

Notes:

1. The source volume ID must be specified when you specify the **-at** *src* parameter.
2. The target volume ID must be specified when you specify the **-at** *tgt* parameter.

SourceVolumeID:TargetVolumeID ... | -
(Required). Specifies that a Remote Mirror and Copy volume relationship for the source and target volume pairs with the specified IDs are to be paused.

Note: Provide a single volume ID instead of a volume pair if you use the **-unconditional** parameter. Specifying pairs results in a format error.

This parameter accepts fully qualified volume IDs, which includes storage image IDs, or a shortened version without storage image IDs if the **-dev** parameter is specified. You must separate multiple remote mirror and copy pair IDs with spaces.

A Remote Mirror and Copy pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a Remote Mirror and Copy relationship. You must separate the two volume IDs of a remote mirror and copy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Note: Requests directed to target volumes are sent to the Storage Image identified by the **-dev** parameter or is embedded in the fully qualified single volume IDs.

Example of a fully qualified pair: IBM.2107-75FA120/0100:IBM.2107-75FA150/0100

Example of a shortened version: 0100:0100

Example of multiple pairs: 0101:0101 0102:0102 0103:0103

Example

Invoking the pausepprc command

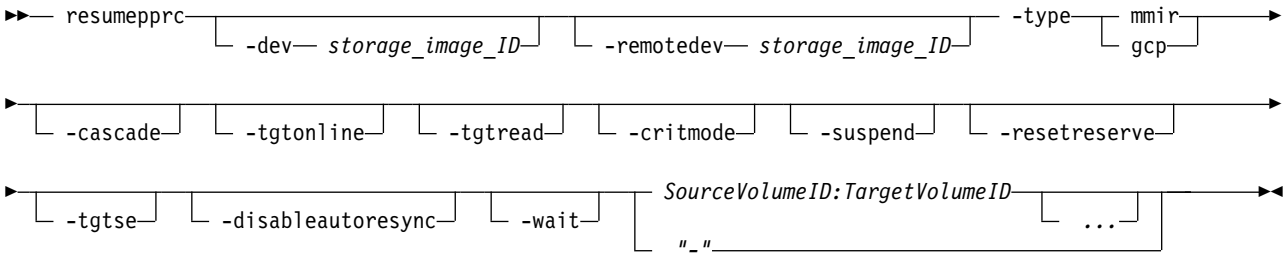
```
dscli> pausepprc -dev IBM.2107-75FA120
-remotedev IBM.2107-75FA150 0100:0100
```

The resulting output

```
Remote Mirror and Copy pair IBM.2107-75FA120/0100:0103
 successfully paused.
```

resumepprc

The **resumepprc** command resumes a remote mirror and copy (formerly PPRC) relationship for a volume pair.



Parameters

Notes:

1. When you specify subsystem IDs, the source and target volumes are restricted to one LSS for the source and one LSS for the target.
2. When you use the **-wait** parameter, periodically issue the **lspprc** command. This command allows you to verify which of the states that the pair has reached during processing.
3. You cannot issue other commands while the **-wait** parameter is processing. The entire transaction must complete before you can proceed with commands other than status commands like **list** commands or **show** commands.

-dev *storage_image_ID*
(Optional) Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified source volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-remotedev *storage_image_ID*
(Optional) Specifies the target storage image ID, which includes manufacturer, machine type, and serial number. This parameter is required if you do not specify a fully qualified target volume ID or if you specify the **-dev** parameter.

-type mmir | gcp
(Required) Modifies one or more remote mirror and copy volume relationships to be either Metro Mirror (synchronous) or Global Copy (extended distance) relationships.

mmir Metro Mirror processing maintains the remote mirror and copy relationship in a consistent (synchronous) manner when the updates are committed to the target. This process becomes slower as the physical distance between source and target increases.

gcp Global Copy processing maintains the remote mirror and copy relationship in a nonsynchronous manner when the updates are committed to the source. Updates to the target volume are performed at a later point in time. The original order of updates is not strictly maintained.

-cascade
(Optional) Enables a remote mirror and copy target volume to be a remote mirror and copy source volume for a different remote mirror and copy volume relationship.

-tgtonline
(Optional) Establishes a remote mirror and copy volume relationship, including when the target volume is online to host systems. This parameter applies to IBM Z volumes. It does not apply to open systems volumes.

-tgtread
(Optional) Allows host servers to read from the remote mirror and copy target volume. For a host server to read the volume, the remote mirror and copy pair must be in a full-duplex state. This parameter applies to open systems volumes and does not apply to IBM Z volumes. The default value for this parameter is disabled.

-critmode
(Optional) Protects the source volume from receiving new data. If the last path fails between the pairs, which results in the inability to send information to the target, the source is protected. Current updates and subsequent attempts to update the source fail, with a unit check on IBM Z.

-suspend
(Optional) Specifies that the remote mirror and copy relationship be suspended when the task completes. This parameter is not valid for a Global Copy (extended distance) remote mirror and copy volume relationship. This parameter is not valid for a Metro Mirror (synchronous) remote mirror and copy volume relationship that is established with the No Copy option.

-resetreserve
(Optional) Establishes the remote mirror and copy relationship when the volume on the secondary logical subsystem is reserved by another host. If this parameter is not specified and the volume on the secondary logical subsystem is reserved, the command fails.

This parameter can only be used with fixed block volumes.

-tgtse
(Optional) Specifies that the secondary volume might be a space-efficient logical volume.

-disableautoresync
(Optional) Allows you to disable the mechanism that automatically resumes a suspended Global Copy relationship. The default is not disabled.

-wait
(Optional). Specifies that the command response be delayed until the remote copy and mirror volumes are in one of the final states: simplex, full duplex, suspended, secondary full duplex,

secondary suspended or secondary simplex (until the pair is not in the Copy Pending state). This parameter cannot be used with the **-type** *gcp* or **-mode** *nocp* parameters.

SourceVolumeID:TargetVolumeID ... | -
(Required) Specifies that a remote mirror and copy volume relationship for the source and target volume pairs with the specified IDs be resumed.

This parameter accepts fully qualified volume IDs, which includes storage image IDs or a shortened version without storage image IDs, if the **-dev** parameter is specified. You must separate multiple remote mirror and copy pair IDs with spaces.

A remote mirror and copy pair ID consists of two volume IDs: one designated as the source and the other as the target volume for a remote mirror and copy relationship. You must separate the two volume IDs of a remote mirror and copy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the resumepprc command

```
dscli> resumepprc -dev IBM.2107-75FA120
-remotedev IBM.2107-75FA150 0100:0100
```

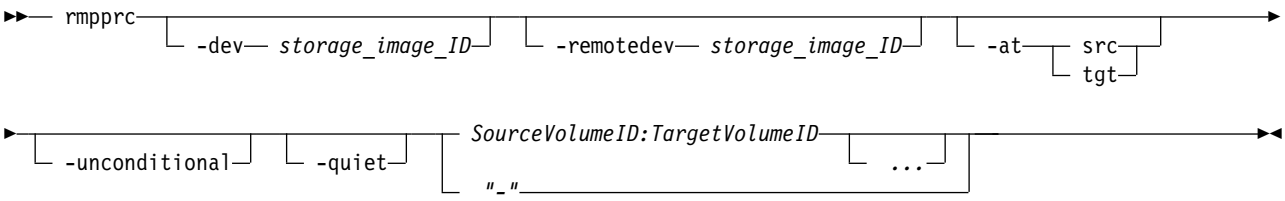
The resulting output

```
Remote Mirror and Copy volume pair IBM.2107-75FA120/0100:0103
 relationship successfully resumed. This message is being returned
 before the copy completes.
```

rmpprc

The **rmpprc** command removes a Remote Mirror and Copy volume pair relationship or removes a single volume ID.

You can remove a single volume ID when a disaster occurs, which allows you to specify only the available volume and not both the primary and secondary volumes. You must specify either the **-at** *src* parameter option or the **-at** *tgt* parameter option when you process a single volume. If neither of these options are specified in the **rmpprc** command, a single volume cannot be processed. The **-unconditional** parameter must also be specified when you process a single volume; otherwise, an error occurs and the command process fails.



Parameters

Notes:

1. When you specify subsystem IDs, the source and target volumes are restricted to one LSS for the source and one LSS for the target.
2. If there is a communication problem between the primary server and the secondary server (two-site configuration) when the **rmpprc** command is issued, the following actions occur:
   - Error message CMUN03012E is issued. This error message indicates that there was a communication problem between the primary and secondary server and that the transaction has failed. However, a partial removal of the pair relationship has occurred.
   - The pair relationship is ended on either the primary volumes or the secondary volumes and the volumes that had the relationship removed enter a simplex state. If a volume is in a simplex state, the volume is not in a Copy Services relationship.

   If this circumstance has occurred, reissue the **rmpprc** command using the **-at** *src* or the **-at** *tgt* parameter and the **-unconditional** parameter for each volume that does not have its pair relationship removed.

   The following list represents the format of the **rmpprc** command when you must remove a partial pair relationship:
   - If the source volume has not been removed from the pair relationship, enter the **rmpprc** command at the **dscli** command prompt with the following parameters and variables:

     ```
     dscli> rmpprc -dev storage_image_ID -at src
     -unconditional SourceVolumeID
     ```
   - If the target volume has not been removed from the pair relationship, enter the **rmpprc** command at the **dscli** command prompt with the following parameters and variables:

     ```
     dscli> rmpprc -dev storage_image_ID -at tgt
     -unconditional TargetVolumeID
     ```

     The value of the storage image ID must be the secondary server.

     The management console must be able to communicate with the secondary server for this command to process successfully.
3. If a disaster occurs involving a three-site configuration, the **rmpprc** command with the **-at** *tgt* and **-unconditional** parameters are used in the recovery process.

-dev *storage_image_ID*
(Optional). Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for all source volumes, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-remotedev *storage_image_ID*
(Optional most times). Specifies the target storage image ID, which includes manufacturer, machine type, and serial number.

Note: The **-remotedev** parameter is required when volume pairs are specified and the **-dev** parameter is specified, as well.

-at src | tgt
(Optional). Specifies that you want to initiate a break action from either a source volume or a target volume.

src Select the **-at** *src* parameter option to initiate a break action from the source volume. After the task successfully runs, the source and target volumes are in the **simplex** state.

tgt Select the **-at** *tgt* parameter option to initiate a break action from the target volume. After the

command successfully runs, the target volume is in the **simplex** state, but there is no guarantee that the source volume state will change. For a break action issued to the target, the source can remain in the **suspend** state.

The **-at** *tgt* parameter option can also be used with single volumes but you must also specify the **-unconditional** parameter. When you specify a single volume using this parameter, the volume is treated as a target and the source is treated as a null.

Notes:

1. If you specify the **-at** *tgt* or **-at** *src* parameter and the **-unconditional** parameter, the value for the **-remotedev** parameter is ignored.
2. If you specify the **-at** *tgt* parameter and do not specify the **-unconditional** parameter, the values for the **-dev** and *SourceVolumeID* parameters are ignored.
3. If you specify the **-at** *src* parameter and do not specify the **-unconditional** parameter, the values for the **-remotedev** and *TargetVolumeID* parameters are ignored.

-unconditional
(Optional). Specifies that a source or target volume has been selected individually, and not as a pair. The **-unconditional** parameter is valid only if the **-at** parameter is selected. Do not use volume pair IDs.

Notes:

1. The source volume ID must be specified when you specify the **-at** *src* parameter.
2. The target volume ID must be specified when you specify the **-at** *tgt* parameter.

-quiet
(Optional). Turns off the Remote Mirror and Copy relationship removal confirmation prompt for this command.

SourceVolumeID:TargetVolumeID ... | -
(Required). Specifies the Remote Mirror and Copy volume relationship for the source and target volume pairs that is to be removed.

Note: Provide a single volume ID instead of a volume pair if you use the **-unconditional** parameter. Specifying pairs results in a format error.

This parameter accepts fully qualified volume IDs, which includes storage image IDs or a shortened version without storage image IDs, if the **-dev** parameter is specified. You must separate multiple Remote Mirror and Copy pair IDs with spaces.

A Remote Mirror and Copy pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a Remote Mirror and Copy relationship. You must separate the two volume IDs of a Remote Mirror and Copy pair ID with a colon and no space. The first volume ID is the source volume. The second volume ID is the target volume.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmpprc command

```
dscli> rmpprc -dev IBM.2107-75FA120
-remotedev IBM.2107-75FA150 0100:0100
```

The resulting output

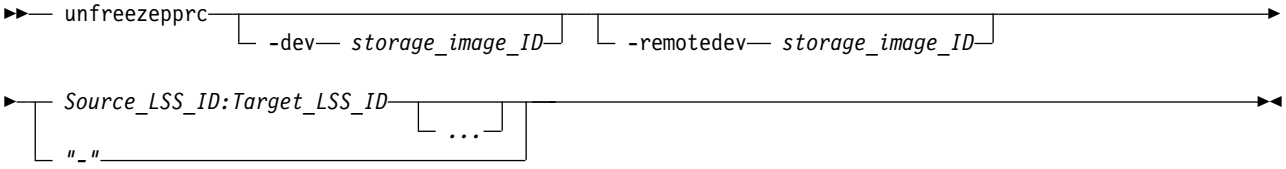
```
Are you sure you want to delete the Remote Mirror and Copy
volume pair relationship 0100:0100? [y/n]: Y

Remote Mirror and Copy pair IBM.2107-75FA120/0100:0100
 successfully removed.
```

unfreezepprc

The **unfreezepprc** command resumes I/O activity on a storage system where the **freezepprc** command has been issued.

The **unfreezepprc** command resets the *queue full* condition for the primary volume. All queued writes to the source volume are written.



Parameters

Notes:

1. This command affects all remote copy and mirror primary volumes that are contained by the LSS(s) that are defined by the Source_LSS_ID:Target_LSS_ID source volume.
2. When specifying subsystem IDs, the command is limited to one LSS pair.
3. Resuming I/O activity on a storage system where the **freezepprc** command has been issued is sometimes referred to as the *thaw* operation.

-dev *storage_image_ID*
(Optional). Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-remotedev *storage_image_ID*
(Optional). Specifies the target storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified target LSS ID or if the **-dev** parameter is used.

Source_LSS_ID:Target_LSS_ID ... | -
(Required). Specifies that a consistency group for the source and target LSSs with the IDs specified be removed from the long busy state.

This parameter accepts fully qualified LSS IDs, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified.

A remote mirror and copy path LSS pair ID consists of two LSS IDs, one designated as the source LSS and the other as the target LSS for a remote mirror and copy path relationship. The two LSS IDs must be separated with a colon and no spaces. Multiple remote mirror and copy path LSS pair IDs must be separated with a space between each value. The first LSS ID is the source LSS. The second LSS ID is the target LSS.

The fully qualified LSS ID format is ***storage_image_ID/xx***, where *xx* is a hexadecimal number in the range 00 - FE.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example pair: 00:00

Example of multiple pairs: 00:00 01:01 02:02

Example

Invoking the unfreezepprc command

```
dscli> unfreezepprc -dev IBM.2107-75FA120
-remotedev IBM.2107-75FA150 01:01
```

The resulting output

```
Remote Mirror and Copy pair ID 01:01 successfully thawed.
```

Global Mirror commands

Global Mirror commands are used to create, manage, view, and delete Global Mirror relationships.

The following Global Mirror commands are available:

mkgmir Starts Global Mirror processing for a specified session.

pausegmir
Pauses Global Mirror processing for the specified session. There are 2 primary reasons to pause Global Mirror processing:

- To repair a part of the Global Mirror infrastructure, such as: Global Copy volume pairs or FlashCopy pairs
- To make modifications to the Global Mirror tuning parameters

resumegmir
Resumes Global Mirror processing for a specified session. If you have issued a **pausegmir** command to pause Global Mirror processing, issue the **resumegmir** command to continue Global Mirror processing.

rmgmir Ends Global Mirror processing for the specified session.

showgmir
Generates two reports. The first report displays the detailed properties about the current Global Mirror operations that are associated with a specified logical subsystem ID. The second report displays the performance metrics for the current Global Mirror operations that are associated with a specified logical subsystem ID.

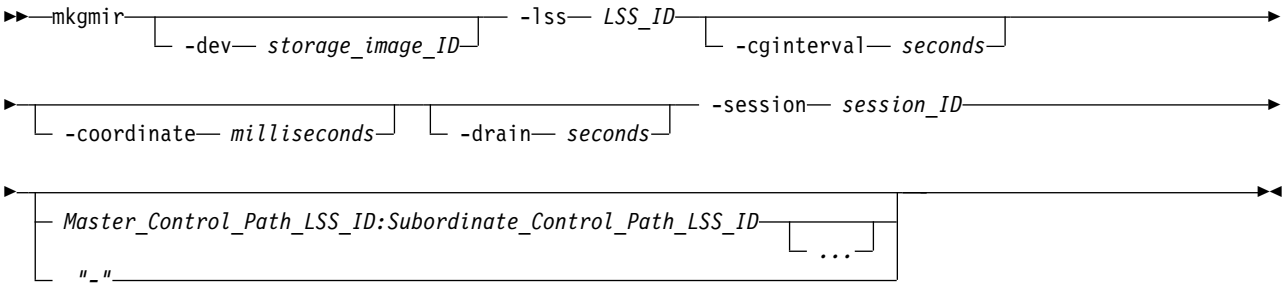
lsgmir Displays a list of Global Mirror for the storage image of the specified logical subsystem (LSS).

showgmircg
Generates a report that displays the consistency group status for the specified Global Mirror session.

showgmiroos
Generates a report that displays the number of unsynchronized (out of sync) tracks for the specified Global Mirror session.

mkgmir

The **mkgmir** command starts Global Mirror for a specified session.



Parameters

Note: If you are using the Cisco MDS 9216 Multilayer Fabric Switch, you must not enable the write acceleration feature. The **mkgmir** command might fail if the write acceleration feature is enabled.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID. It is also required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. The **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-lss *LSS_ID*
(Required) Specifies the master logical subsystem (LSS) that receives the **mkgmir** command. This parameter accepts a fully qualified master LSS ID. The parameter includes either the storage image ID or a shortened version without the storage image ID if you specify the **-dev** parameter.

-cginterval *seconds*
(Optional) Specifies the consistency group interval time, in seconds. This number specifies how long to wait between the formation of consistency groups. If this number is not specified or is set to zero, consistency groups are formed continuously. The consistency group interval setting is required for a start action. If not specified, the default interval is zero. The consistency group interval setting can be modified for a resume action; otherwise, the interval that is specified for the start action is maintained.

The maximum value is 65,535 seconds.

-coordinate *milliseconds*
(Optional) Specifies the maximum coordination interval, in milliseconds. This value indicates the maximum time that Global Mirror processing queues Primary/Host/IO to start forming a consistency group. The coordination interval setting is required for a start action. If this value is not specified, the default interval is 50 milliseconds.

The coordination interval setting can be modified for a resume action: otherwise, the interval that is specified for the start action is maintained. The maximum value is 65,535 milliseconds.

-drain *seconds*
(Optional) Specifies the maximum consistency group drain time in seconds and the maximum amount of time that the consistent set of data is allowed to drain to the remote site before the current consistency group fails. The drain time setting is required for a start action. If the drain time is not specified, the default drain time is 240 seconds.

The drain time setting can be modified for a resume action; otherwise, the time that is specified for the start action is maintained.

-session *session_ID*
(Required) Specifies that Global Mirror for the specified session is start or resume. A session ID is a Global Mirror session number that you assign in the 01 - FF hexadecimal range.

Example: 01

Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID ... | -
(Optional) Specifies one or more Global Mirror associations. A Global Mirror association consists of two fully qualified LSS IDs. The first is designated as the master resource and the second is the subordinate resource between a PPRC path that was established. An LSS ID is a two-character value hexadecimal value in the range 00 - FE. You must separate the fully qualified LSS IDs of a Global Mirror association with a colon and no spaces. The master resource must be identical for all relationships.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive mode.

Example

Invoking the mkgmir command

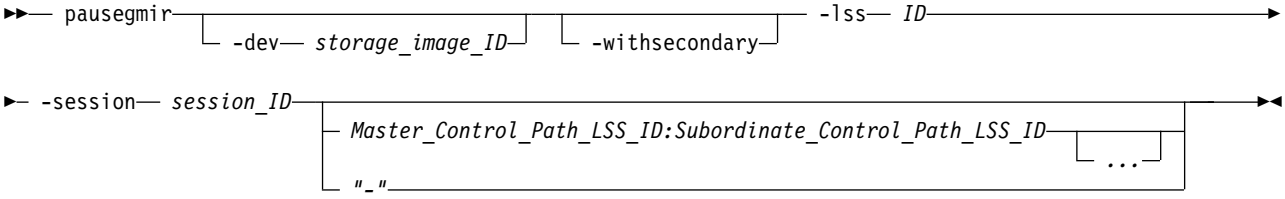
```
dscli> mkgmir
-dev IBM.2107-75FA120
-lss 10 -session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00
```

The resulting output

```
Global Mirror for session 01 successfully started.
```

pausegmir

The **pausegmir** command pauses Global Mirror for the specified session.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-withsecondary
(Optional) Specifies that once the Global Mirror session is paused, the secondary volumes will form a consistent copy of the data set.

-lss *ID*
(Required) Specifies the master logical subsystem (LSS) that receives the **pausegmir** command.

Accepts a fully qualified master LSS ID, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified.

-session *session_ID*
(Required) Specifies the session ID for which the Global Mirror process is to be paused. A session ID is a Global Mirror session number that you assign in the 01 - FF hexadecimal range.

Example: 01

Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID ... | **-**
(Optional). Specifies one or more Global Mirror path associations. A Global Mirror (Asynchronous PPRC) path association consists of two fully qualified LSS IDs, one designated as the master resource and the other as the subordinate resource between which a remote copy and mirror path has been established.

A LSS ID is two hexadecimal characters in the range 00 - FE. You must separate the fully qualified LSS IDs of a Global Mirror path association with a colon and no spaces. The master resource must be identical for all relationships.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive mode.

Example of one Global Mirror association with a single subordinate in the configuration: IBM.2107-75FA120/00:IBM.2107-75FA150/00

Example of multiple Global Mirror associations with two subordinates in the configuration: IBM.2107-75FA120/00:IBM.2107-75FA150/00 IBM.2107-75FA120/00:IBM.2107-75FA150/01

Example

Invoking the pausegmir command

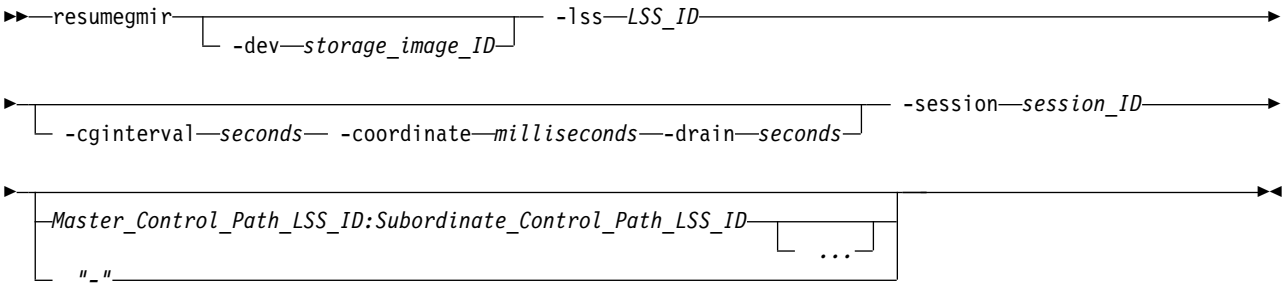
```
dscli> pausegmir
-dev IBM.2107-75FA120 -lss 10
-session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00
```

The resulting output

```
Global Mirror for session 01 successfully paused.
```

resumegmir

The **resumegmir** command resumes Global Mirror processing for a specified session.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set

the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-lss *LSS_ID*
(Required). Specifies the master logical subsystem (LSS) that is to receive the **resumegmir** command. Accepts a fully qualified LSS ID, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is two hexadecimal digits in the range 00 - FE.

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

Tuning parameters consist of the following three values: -cginterval *seconds*, **-coordinate** *milliseconds*, **-drain** *seconds*
Tuning parameters have default values applied to them from the microcode. However, you can choose to change those values. You must designate a value for each of the parameters even if you are changing only one value. For example, if you decide to change only the value on the **-cginterval** parameter from 0 to 1, your command must include the default values for the other two parameters.

```
dscli> resumegmir IBM.2107-75FA120/10 -cginterval 1
-coordinate 50 -drain 30 -session 01
IBM.2107-75FA120/00:IBM.2107-75FA150/00
```

-cginterval *seconds*: Specifies the consistency group interval time, in seconds. This number specifies how long to wait between the formation of consistency groups. If this number is not specified or is set to zero, consistency groups are formed continuously.
The default value is 0. The maximum value is 65 535 seconds.

-coordinate *milliseconds*: Specifies the maximum coordination interval, in milliseconds. This value indicates the maximum time that Global Mirror processing queues Primary/Host/IO to start forming a consistency group.
The default value is 50 milliseconds. The maximum value is 65 535 milliseconds.

-drain *seconds*: Specifies the maximum consistency group drain time in seconds and the maximum amount of time that the consistent set of data is allowed to drain to the remote site before failing the current consistency group.
The default drain time is 30 seconds.

-session *session_ID*
(Required) Specifies the Global Mirror session that is to be started. A session ID is a Global Mirror session number that you assign in the 01 - FF hexadecimal range.

Example: 01

Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID ... | **-**
(Optional). Specifies one or more Global Mirror path associations. A Global Mirror path association consists of two fully qualified LSS IDs. The first is designated as the master resource and the second is the subordinate resource between which a PPRC path has been established. A LSS ID is two hexadecimal characters in the range 00 - FE. You must separate the fully qualified LSS IDs of a Global Mirror association with a colon and no spaces. The master resource must be identical for all relationships.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive mode.

Example of one Global Mirror association with a single subordinate in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00

Example of multiple Global Mirror associations with two subordinates in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00 IBM.2107-75FA120/00: IBM.2107-75FA150/01

Example

Invoking the resumegmir command

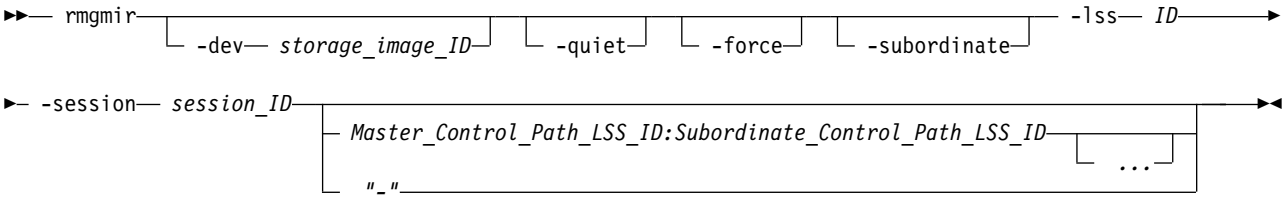
```
dscli> resumegmir
-dev IBM.2107-75FA120 -lss 10
-session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00
```

The resulting output

```
Global Mirror for session 01 successfully resumed.
```

rmgmir

The **rmgmir** command ends Global Mirror processing for the specified session.



Parameters

Note: Although this command might interrupt the formation of a consistency group, every attempt is made to preserve the previous consistent copy of the data on the FlashCopy target volumes. If, due to failures, this command cannot complete without compromising the consistent copy, the command stops processing and an error code is issued. If this occurs, reissue this command (**rmgmir**) with the **-force** parameter to force the command to stop the Global Mirror process.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-quiet
(Optional) Turns off the end Global Mirror session confirmation prompt for this command.

-force
(Optional) Forces the Global Mirror process to stop regardless of the state of the Global Mirror associations.

-subordinate
(Optional) Indicates that the **-lss** parameter specifies a subordinate LSS ID.

-lss *ID*
(Required) Specifies the logical subsystem (LSS) that is participating in the Global Mirror session. Accepts a fully qualified LSS ID, which includes the storage image ID or a shortened version without the storage image ID, if the **-dev** parameter is specified.

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

-session *session_ID*
(Required) Specifies the session ID for which the Global Mirror path association will be removed. A session ID is a Global Mirror session number that you assign in the 01 - FF hexadecimal range.

Example: 01

Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID ... | **-**
(Optional) Specifies one or more Global Mirror path associations. A Global Mirror path association consists of two fully qualified LSS IDs. The first is designated as the master resource and the second is the subordinate resource between which there is a remote mirror and copy path. A LSS ID is two hexadecimal characters in the range 00 - FE. You must separate the fully qualified LSS IDs of a Global Mirror association with a colon and no spaces. The master resource must be identical for all relationships.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive mode.

Example of one Global Mirror path association with a single subordinate in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00

Example of multiple Global Mirror path associations with two subordinates in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00 IBM.2107-75FA120/00: IBM.2107-75FA150/01

Example

Invoking the rmgmir command

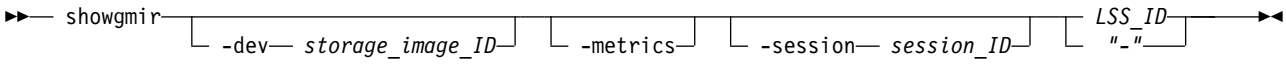
```
dscli> rmgmir
-dev IBM.2107-75FA120 -lss 10
-session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00
```

The resulting output

```
Are you sure you want to stop Session ID 01? [y/n]: Y
Global Mirror for Session ID 01 successfully stopped.
```

showgmir

The **showgmir** command displays properties and performance metrics for a Global Mirror logical subsystem ID. You can issue this command on either the master storage system or on any of the subordinate storage systems. The report that is generated by this command varies significantly depending on which storage system that you issue the command and the parameters that you specify.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-metrics
(Optional) Specifies that the logical subsystem ID and its performance statistics be displayed.

-session *session_ID*
(Optional) Specifies a session ID number. The number must be greater than 0 (01 - FF hexadecimal range). If you do not specify this parameter, all sessions will be displayed.

LSS_ID | **-**
(Required) Specifies the logical subsystem (LSS) that receives the **showgmir** command. This parameter accepts a fully qualified LSS ID, which includes the storage image ID or a shortened version without the storage image ID if the **-dev** parameter is specified. The LSS ID is two hexadecimal digits in the range 00 - FE.

If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive mode.

Note:

The type of report that you receive is determined by the value that you specify for the LSS ID as follows:

- When you issue the **showgmir** command from the master storage system and you specify an LSS ID that is the same type (both even numbers or both odd numbers) as the master, you receive the following results:
  - Without the **-metrics** parameter: A properties report that includes the master information
  - With the **-metrics** parameter: A properties and performance values report.
- When you issue the **showgmir** command from the master storage system and you specify an LSS ID that is not the same type (one even number and one odd number) as the master, you receive the following results:
  - Without the **-metrics** parameter: A properties report that displays only the fully qualified LSS_ID value and all the other input fields display a "-" value.
  - With the **-metrics** parameter: A properties and performance report that displays only the fully qualified LSS_ID value and all the other input fields display a "-" value.
- When you issue the **showgmir** command from the subordinate storage system and you specify an LSS ID that is the same type (both even numbers or both odd numbers) as the master, you receive the following results:
  - Without the **-metrics** parameter: A detailed properties report that displays only the master information (fully-qualified LSS ID, master session ID, and master storage system ID). All the other fields display as a "-" value.
  - With the **-metrics** parameter: A detailed properties and performance values report that displays only the master information (fully qualified LSS ID, master session ID, and master storage system ID). All the other fields display as a "-" value.
- When you issue the **showgmir** command from the subordinate storage system and you specify an LSS ID that is *not* the same type (one even number and one odd number) as the master, you receive the following results:
  - Without the **-metrics** parameter: A properties report that displays only the fully qualified LSS_ID value. All the other input fields on the report display a "-" value.
  - With the **-metrics** parameter: A properties and performance report that displays only the fully qualified LSS_ID value. All the other input fields on the report display a "-" value.

Example (metrics not specified)

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **showgmir** command.

Note: The following example presumes that you have issued the **showgmir** command with an LSS ID that matches the master storage system.

Invoking the showgmir command

```
dscli> showgmir -dev IBM.2107-1300861 14
```

The resulting output

ID	Master Count	Master Session ID	Copy State	Fatal Reason	CG Interval Time (seconds)	Coord. Time (milli-seconds)	Max CG Drain Time (seconds)	Current Time
IBM.2107-1300861/14	1	0x25	Running	Not Fatal	0	50	30	04/25/2017 13:45:44 PDT

CG Time	Successful CG Percentage	FlashCopy Sequence Number	Master ID	Subordinate Count	Master/ Subordinate Assoc
04/25/2017 13:45:44 PDT	100	0x453-FCCF8	IBM.2107-1300861	1	IBM.2107-1300861/14: IBM.2107-1300321/14

Report field definitions

ID Indicates the LSS ID that consists of a storage image ID followed by two hexadecimal characters that identify a Global Mirror (asynchronous remote mirror and copy) master LSS ID. The two hexadecimal digits is in the range 00 - FE.

Master Count
Indicates the number of Global Mirror (asynchronous remote mirror and copy) masters. This value could be zero if none exist. When the value is zero, the master information fields of the report display a "-" value

Master Session ID
Indicates the session ID that you assigned, 01 - FF hexadecimal range.

Copy State
Indicates the Global Mirror (asynchronous remote mirror and copy) copy state. One of the following values is displayed:

`Running`
Indicates that the Global Mirror copy process is running.

`Paused`
Indicates that the Global Mirror session will be paused, that is, stopped from forming consistency groups, after the current consistency group has been formed. However, the secondary volumes in the Global Mirror session should not be considered to form a consistent data set. Another separate process is required to form a consistent data set on the secondary volumes. You can pause a session and later resume the session.

`Pause In Progress`
Indicates that the Global Mirror copy process is in the process of pausing.

`Paused with Secondary Consistency`
Indicates that the Global Mirror session will be paused, that is, stopped from forming of consistency groups, after the current consistency group has been formed. However, all of the secondary volumes in the Global Mirror session already form a consistent data set. You can pause a session and later resume the session.

`Paused because Resume Failed`
Indicates that an attempt to resume a Global Mirror session that is in the *Paused* state failed. The Global Mirror session is still in the *Paused* state.

Fatal
Indicates that the Global Mirror copy process is failed.

Unowned
Indicates that the session is not owned by the cluster that you specified using the LSS ID.

Recovering
Indicates that the Master is in the process of recovering the session.

Fatal Reason
Indicates a reason code for the failure. One of the following values is displayed:

Time out

Revert FLC Failed Timeout

Revert FLC Failed

Not Fatal

Invalid Session ID

Inaccessible or Failed

Consistency Check Not Completed

Consistency Check Failed

Communication Failure

CG Corrupted

Busy Condition Preventing

CG Interval Time (seconds)
Indicates the consistency group interval time between attempts to form a consistency group, up to 65 535 seconds.

Coord. Time (milliseconds)
Indicates the maximum extended distance coordination interval. The default time is 50 milliseconds.

Max CG Drain Time (seconds)
Indicates the consistency group drain time. The consistency group drain time is the maximum time that the consistent set of data is allowed to drain to the remote site before failing the current consistency group. The maximum allowed time is 65 535 seconds.

Current Time
Indicates the time stamp for when this report was generated. The date is displayed in the *MM/DD/YYYY* format. The time is displayed in the *HH:MM:SS* format on a 24-hour clock.

Note: If the clock is automatically adjusted at applicable times between standard and daylight saving time, daylight saving time is set to 1. If the clock is not automatically adjusted for daylight saving time, set to 0. For example, the report would display 12/04/2006 08:00:00 MST 0 if the clock is not automatically adjusted for daylight saving time.

CG Time
Indicates the recorded time stamp when the last successful consistency group was formed.

Successful CG Percentage
Indicates the percentage of successful attempts to form a consistency group, from 0% to 100%.

FlashCopy Sequence Number
Indicates the FlashCopy sequence number that is associated with the current consistency group.

Note: This value does not apply to a 2105; a "-" value is displayed in this column when a machine type 2105 is queried.

Master ID
Indicates the Global Mirror master storage image ID.

Subordinate Count
Indicates the count of subordinate associations (with an allowed value of 1 to 16). The master-subordinate association fields repeat according to this count.

Master/Subordinate Assoc
Indicates the Global Mirror path associations. A Global Mirror path association consists of two fully qualified LSS IDs. One ID is designated as the master resource and the other ID is designated as the subordinate resource; a remote mirror and copy path has been established between the two resources.

Example (metrics specified)

The following tables represent the headers that are displayed on the output report that is associated with the **showgmir** command using the **-metrics** parameter.

Note: The following example presumes that you have issued the **showgmir** command with an LSS ID that matches the master storage system.

Invoking the showgmir command

```
dscli> showgmir -dev IBM.2107-75FA120  -metrics 10
```

The resulting output

ID	Total Failed CG Count	Total Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milli-seconds)	CG Interval Time (seconds)
IBM. 2107-75FA120/10	2	3	40	2	04/13/2017 13:08:37 PST	50	5

Max CG Drain Time (seconds)	First Failure Control Unit	First Failure LSS	First Failure Status	First Failure Reason	First Failure Master State	Last Failure Control Unit	Last Failure LSS
240	IBM.2107-75FA120	0x05	Error	Long Busy	Error Recovery	IBM.2107-75FA120	0x05

Last Failure Status	Last Failure Reason	Last Failure Master State	Previous Failure Control Unit	Previous Failure LSS	Previous Failure Status	Previous Failure Reason	Previous Failure Master State
Error	Long Busy	Error Recovery	IBM.2107-75FA120	0x05	Error	Long Busy	Error Recovery

Report field definitions

ID Indicates the LSS ID. This ID consists of a storage image ID that is followed by two hexadecimal characters that identify a Global Mirror (Asynchronous PPRC) master LSS ID. The two hexadecimal digits is in the range 00 - FE.

Total Failed CG Count
Indicates the total number of consistency groups that did not complete in the user-specified drain time.

Total Successful CG Count
Indicates the total number of consistency groups that completed before the user-specified drain time.

Successful CG Percentage
Indicates the percentage of attempts that were successful in forming a consistency group.

Failed CG after Last Success
Indicates the total number of failed consistency groups after the last successful completion.

Last Successful CG Form Time
Indicates the last successful consistency group completion time.

Coord. Time (milliseconds)
Indicates the value in milliseconds that indicates the maximum amount of time that Global Mirror queues the primary host I/O to start forming a consistency group. The default is 50 milliseconds.

CG Interval Time (seconds)
Indicates the value in seconds that indicates how long to wait between formation of consistency groups.

Max CG Drain Time (seconds)
Indicates the value in seconds that indicates the maximum amount of time that Global Mirror allows for the consistent set of data to drain to the remote site.

First Failure Control Unit
Indicates the Control Unit MTS that has caused the first failure of the consistency group formation.

First Failure LSS
Indicates the LSS number that has caused the first failure of the consistency group formation.

First Failure Status
Indicates the first failure status of the consistency group formation. The “First Failure Reason” and “First Failure Master State” fields display data only when this field contains “Error”.

First Failure Reason
Indicates the error reason of the first failure of the consistency group formation attempt.

First Failure Master State
Indicates the master state for the first Global Mirror failure.

Last Failure Control Unit
Indicates the Control Unit MTS that has caused the last failure of the consistency group formation.

Last Failure LSS
Indicates the LSS number that has caused the last failure of the consistency group formation.

Last Failure Status
Indicates the last failure status of the consistency group formation. The “Last Failure Reason” and “Last Failure Master State” fields display data only when this field contains “Error”.

Last Failure Reason
Indicates the error reason of the last failure of the consistency group formation attempt.

Last Failure Master State
Indicates the master state for the last Global Mirror failure.

Previous Failure Control Unit
Indicates the Control Unit MTS that has caused the previous failure of the consistency group formation.

Previous Failure LSS
Indicates the LSS number that has caused the previous failure of the consistency group formation.

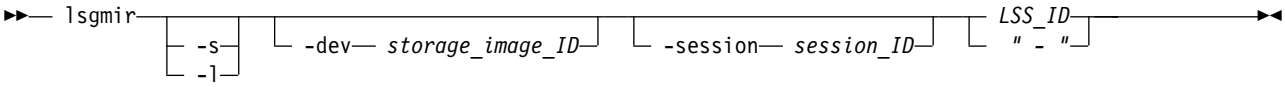
Previous Failure Status
Indicates the previous failure status of the consistency group formation. The "Previous Failure Reason" and "Previous Failure Master State" fields display data only when this field contains "Error".

Previous Failure Reason
Indicates the error reason of the previous failure of the consistency group formation attempt.

Previous Failure Master State
Indicates the master state for the previous Global Mirror failure.

lsgmir

The **lsgmir** command displays a list of Global Mirror for the storage image of the specified logical subsystem (LSS).



Parameters

-s
(Optional) Specifies that you want the system to display only the session ID for the LSS. You cannot use the **-s** and the **-l** parameters together.

-l
(Optional) Specifies that you want the system to display all the information that is associated with the session. You cannot use the **-l** and the **-s** parameters together.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a value for manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified extent pool ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter will temporarily override any defined value for *devid* for the current command.

-session *session_ID*
(Optional) Specifies a session ID number. The number must be greater than 0 (01 - FF hexadecimal range). If you do not specify this parameter, all sessions will be displayed.

LSS_ID | **-**
(Required) Specifies the logical subsystem (LSS) ID for the Global Mirror session. A fully qualified LSS ID is accepted, which consists of the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is a four-decimal digit number with no leading zeros, prefixed with the letter P.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

Invoking the `lsgmir` command to display a list of Global Mirror for the storage image of the specified logical subsystem.

```
dscli> lsgmir -dev IBM.2107-75FA120
```

The resulting output

SessionID	MasterID	ID	State	%Success	CGtime
1	IBM.2107-75FA120	10	Running	90	08/31/2018 21:52:02 CST

Report field definitions

SessionID
Identifies the session ID number.

MasterID
Identifies the master ID.

ID Identifies the logical subsystem ID.

State One of the following Global Mirror states are displayed:

`Running`
Indicates that a Global Mirror session is resumed.

`Paused`
Indicates that the Global Mirror session will be paused, that is, stopped from forming consistency groups, after the current consistency group has been formed. However, the secondary volumes in the Global Mirror session should not be considered to form a consistent data set. Another separate process is required to form a consistent data set on the secondary volumes. You can pause a session and later resume the session.

`Pause in Progress`
Indicates that the Global Mirror session is currently in the process of pausing.

`Paused with Secondary Consistency`
Indicates that the Global Mirror session will be paused, that is, stopped from forming of consistency groups, after the current consistency group has been formed. However, all of the secondary volumes in the Global Mirror session already form a consistent data set. You can pause a session and later resume the session.

`Paused because Resume Failed`
Indicates that an attempt to resume a Global Mirror session that is in the *Paused* state failed. The Global Mirror session is still in the *Paused* state.

`Fatal`
Indicates that the Global Mirror copy process is failed.

`Unowned`
Indicates that the session is not owned by the cluster that you specified using the LSS ID.

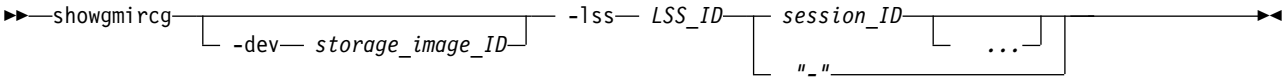
`Recovering`
Indicates that the Master is in the process of recovering the session.

%Success
Indicates the percentage of successful attempts to form a consistency group.

CGtime
Indicates the time stamp of the last successful consistency group.

showgmircg

The **showgmircg** command displays consistency group status for the specified Global Mirror session.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-lss *LSS_ID*
(Required) Specifies the master logical subsystem (LSS) that receives the **showgmircg** command. LSS ID consists of two hexadecimal characters in the range of 00 - FE.

This parameter accepts a fully qualified master LSS ID, which includes either the storage image ID or a shortened version without the storage image ID if the **-dev** parameter is specified.

session_ID ... | **-**
(Required) Specifies one session to display. A session ID is a Global Mirror session number that you assign in the 01 - FF hexadecimal range.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive mode.

Example: 01

Example

Invoking the showgmircg command

```
dscli> showgmircg -dev IBM.2107-75FA120 -lss 10 01
```

The resulting output

```
LSS ID    IBM.2107-75FA120/10
Session   01
CG Status 0
```

Report field definitions

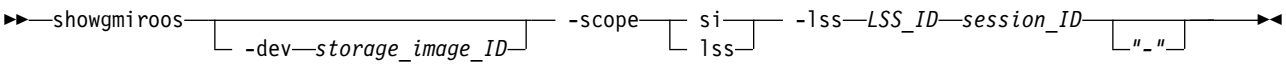
LSS ID
Indicates the logical subsystem ID.

Session
Indicates the Global Mirror session number.

CG Status
Indicates the Global Mirror Consistency Group status (primarily used by Field Engineering).

showgmiroos

The **showgmiroos** command displays the number of unsynchronized (out of sync) tracks for the specified Global Mirror session.



Parameters

Note: You might want to use this command to assist you in the following circumstances:

- To see how data is transferring. The **showgmiroos** command lets you see how far behind the remote site is from the local site in the transfer of data. The displayed value represents how many tracks must be transferred to catch up (be synchronized).
- You are not able to form consistency groups because you have exceeded the maximum drain time. The number of tracks that are not synchronized might be an indicator that you must adjust some values to allow for complete processing.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-scope si | lss
(Required) Specifies the scope of the data to be returned: storage image (si) or logical subsystem (lss).

-lss *LSS_ID*
(Required) Specifies the master logical subsystem (LSS) that receives the **showgmiroos** command. Accepts a fully qualified master LSS ID, which includes either the storage image ID or a shortened version without the storage image ID if the **-dev** parameter is specified. The LSS ID is two hexadecimal digits in the range 00 - FE.

session_ID -
(Required) Specifies the session to display. A session ID is a Global Mirror session number that you assign in the 01 - FF hexadecimal range.

If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive mode.

Example

Invoking the showgmiroos command

```
dscli> showgmiroos -dev IBM.2107-75FA120 -scope si -lss 10 01
```

The resulting output

```
Scope           IBM.2107-75FA120
Session         01
OutOfSyncTracks 3
```

Report field definitions

Scope Indicates the scope of the returned information (storage image or logical subsystem).

Session
Indicates the Global Mirror session number.

OutOfSyncTracks
Indicates the number of unsynchronized tracks.

Global Mirror session commands

Global Mirror commands are used to create, modify, view, and delete Global Mirror sessions.

The following Global Mirror session commands are available:

chsession
Allows you to modify a Global Mirror session.

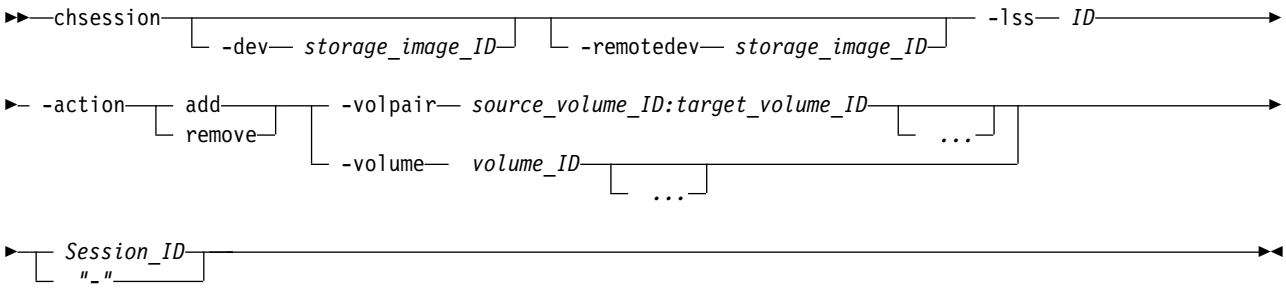
lssession
Generates a report that displays a list of Global Mirror sessions for a logical subsystem (LSS) and information regarding the volumes associated with each session in the list.

mksession
Opens a Global Mirror session.

rmsession
Closes an existing Global Mirror session.

chsession

The **chsession** command allows you to modify a Global Mirror session.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the ID of the storage image containing the logical subsystem. The storage image ID includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID or do not set the *devid* variable in your profile or with the **setenv** command. The storage image ID is also required if the HMC is associated with more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-remotedev *storage_image_ID*
(Optional) Specifies the remote DS system that contains one or more target volume IDs that are defined by the *source_volume_ID:target_volume_ID* variable of the **-volpair** parameter. The remote storage ID includes the manufacturer, machine type, and serial number. This parameter is required if you do not specify a fully qualified target volume ID or if you specify the **-dev** parameter.

-lss *ID*
(Required) The logical subsystem (LSS) ID for the Global Mirror session. The format of the LSS ID is a hexadecimal number in the range 00 - FE.

This parameter accepts a fully qualified LSS ID, which includes either the storage image ID or a shortened version without the storage image ID if the **-dev** parameter is specified.

-action add | remove
(Required). Specify one of the following values.

add Specifies that volumes be added to the session.

remove
Specifies that volumes be removed from the session.

-volpair *source_volume_ID:target_volume_ID ...* | **-volume** *volume_ID ...*
(Required) Specifies an array of one or more volume IDs or volume ID ranges to be added or removed for the Global Mirror session. All volumes must share a common logical subsystem.

The ellipsis (...) indicates that, optionally, you can specify multiple values. To specify a range of volume IDs, you must separate the volume IDs with a hyphen. To specify a combination of one or more volume IDs and a range of volume IDs, separate the volume IDs and ranges with commas.

Notes:

- Specify either the **-volpair** parameter or the **-volume** parameter.
- See the following information for the **-volpair** and **-volume** parameters.

-volpair *source_volume_ID:target_volume_ID ...*
A Remote Mirror and Copy pair ID consists of two volume IDs, one designated as the source and the other as the target for a Remote Mirror and Copy relationship. To specify two volume IDs, separate the volume IDs with a colon and no space. The first volume ID is the designated source volume and the second ID as the designated target volume. The secondary volume ID is required to identify a specific relationship in a Multi-Target Metro Mirror relationship.

To specify fully qualified volume IDs, include the storage image ID or a shortened version without storage image ID if you specify the **-dev** or **-remotedev** parameters.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of 0x*XYZZ*, where:

X Specifies the address group, 0 - F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

Example of multiple volume IDs and ranges: 0200-0204:0400-0404,0210:0410

-volume *volume_ID ...*
The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0 - F.

Example of a range of volume IDs: 0010-001F

Example of multiple volume IDs and ranges: 0100-010F,0180-018F,0120

Session_ID | **-**
(Required) Specifies the Global Mirror session that is modified for this session ID. A session ID is a hexadecimal number in the range 01 - FF.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a session ID: 01

Example

Invoking the chsession command

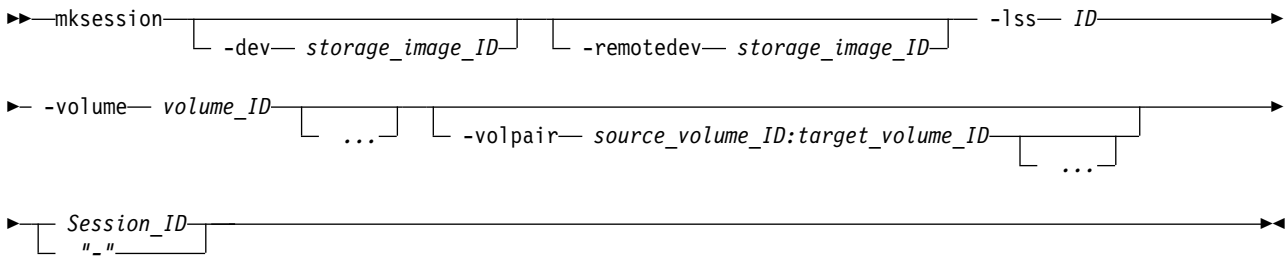
```
dscli> chsession -dev IBM.2107-75FA120 -remotedev IBM.2107-7504491 -lss 10 -action add
-volpair 1000-1010:1000-1010 01
```

The resulting output

```
Global Mirror session 01 successfully modified.
```

mksession

The **mksession** command opens a Global Mirror session.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the logical subsystem or do not set the *devid* variable in your profile or through the **setenv** command. The storage image ID is also required if the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-remotedev *storage_image_ID*
(Optional) Specifies the remote DS system that contains one or more target volume IDs that are defined by the *source_volume_ID:target_volume_ID* variable of the **-volpair** parameter. The remote storage ID includes the manufacturer, machine type, and serial number. This parameter is required if you do not specify a fully qualified target volume ID or if the **-dev** parameter is selected.

-lss *ID*
(Required) Creates a Global Mirror session for this logical subsystem. Accepts a fully qualified LSS ID, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The LSS ID is a hexadecimal number in the range 00 - FE.

-volume *volume_ID ...*
(Required) Specifies an array of one or more volume IDs or a range of volume IDs to be included in the Global Mirror session. All volumes must share a common logical subsystem.

To specify a range of volume IDs, you must separate the volume IDs with a hyphen. To specify a combination of one or more volume IDs and a range of volume IDs, separate the volume IDs and ranges with commas. The ellipsis (...) indicates that, optionally, you can specify multiple values.

Do not qualify the volume ID with the storage image ID. The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

X Specifies the address group, 0 - F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

Example of a volume ID: 0010

Example of a range of volume IDs: 0010-001F

Example of multiple volume IDs and ranges: 0100-010F,0180-018F,0120

-volpair *source_volume_ID:target_volume_ID ...*
(Optional) Specifies an array of one or more volume pair IDs or volume pair ranges to include in the Global Mirror session. All volumes must be of a common LSS ID.

Either the **-volume** or **-volpair** parameter is required but you cannot specify both together.

The ellipsis (...) indicates that, optionally, you can specify multiple values. To specify a range of volume IDs, separate the volume IDs with a hyphen. To specify a combination of one or more volume IDs and a range of volume IDs, separate the volume IDs and ranges with commas.

To specify fully qualified volume IDs, include the storage image ID or a shortened version without storage image ID if you specify the **-dev** or **-remotedev** parameters.

A Remote Mirror and Copy pair ID consists of two volume IDs, one designated as the source and the other as the target volume for a Remote Mirror and Copy relationship. To specify two volume IDs, separate the volume IDs with a colon and no space. The first volume ID is the designated source volume. The second volume ID is the designated target volume. The secondary volume ID is required to identify a specific relationship in a Multi-Target Metro Mirror relationship.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ*, where:

X Specifies the address group, 0 - F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

Example of multiple volume IDs and ranges: 0200-0204:0400-0404, 0100:0300

Session_ID | **-**
(Required) Specifies the session ID for which Global Mirror processing is allowed. A session ID is a hexadecimal number in the range 01 - FF.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of a session ID: 01

Example

Invoking the mksession command

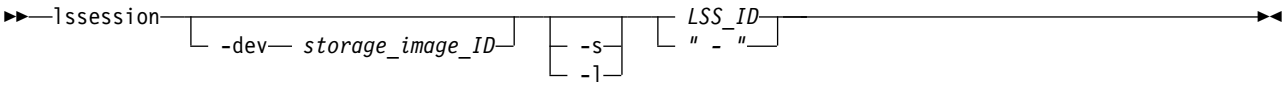
```
dscli> mksession -dev IBM.2107-75FA120 -lss 10 –volume 1000-100F 01
```

The resulting output

```
Global Mirror session ID 01 successfully opened.
```

lssession

The **lssession** command displays a list of Global Mirror sessions for a logical subsystem (LSS) and information about the volumes of each session in the list.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID or do not set

the *devid* variable in your profile or through the **setenv** command. The storage image ID is also required if the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-s
(Optional) Displays the session IDs. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output. You cannot use the **-l** and the **-s** parameters together.

LSS_ID | **-**
(Required) Specifies the logical subsystem (LSS) ID for the Global Mirror session. The format of the LSS ID is a hexadecimal value in the range 00 - FE.

This parameter accepts a fully qualified LSS ID, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified.

If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lssession** command and the **-l** parameter.

When you use the **-s** parameter with the **lssession** command, only three ID items are displayed in the resulting report: LSSID, SessionID, and VolumeID. A separate example is shown for this scenario.

Invoking the lssession command using the -l parameter

```
dscli> lssession –dev IBM.2107-75FA120 -l 01
```

The resulting output

LSSID	Session	Status	Volume	Volume-Status
IBM.2107-75FA120/10	01	Normal	IBM.2107-75FA120/1001	Active
IBM.2107-75FA120/10	01	Normal	IBM.2107-75FA120/1002	Active
IBM.2107-75FA120/10	01	Normal	IBM.2107-75FA120/1003	Active

Primary-Status	Secondary-Status	First-Pass-Complete	Allow-Cascading
Primary Full Duplex	Secondary Simplex	True	Disabled
Primary Full Duplex	Secondary Simplex	True	Disabled
Primary Full Duplex	Secondary Simplex	True	Disabled

Invoking the lssession command using the -s parameter

```
dscli> lssession -s –dev IBM.2107-75FA120 10
```

The resulting output

LSSID	Session	Volume
10	01	1001
10	01	1002
10	01	1003

Report field definitions

LSSID
Indicates the unique identifier that is assigned to this logical subsystem object. The LSS ID is a hexadecimal value in the range 00 - FE.

Session
Indicates the Session ID number that you assigned in the 01 - FF hexadecimal range.

Status Indicates the state of the session. One of the following values is displayed:

CG in progress
Indicates that the consistency group for the session is in progress.

Increment Pending
Indicates that the Increment process is in progress.

Normal
Indicates that the session is in a normal state.

Volume
Indicates the volume ID. If no volume is active for the session a " - " value is displayed.

VolumeStatus
Indicates the status of the volume in the session. One of the following values is displayed:

Join Pending
Indicates that the volume is not active for the current session. However, it is added to the session in the next cycle.

Remove Pending
Indicates that the volume is active for the current session. However, it is removed in the next cycle.

Active Indicates that the volume is an active member of the session.

PrimaryStatus
Indicates the primary remote copy and mirror status of the volume. One of the following values is displayed:

Primary Simplex
Indicates that the volume is not part of a remote mirror and copy relationship.

Primary Copy Pending
Indicates that the volume is primary in a remote mirror and copy relationship and the relationship is in a Copy Pending state, which means that the source and target volume are out-of-sync. In this situation, data still needs to be copied from the source to the target volume.

Primary Full Duplex
Indicates that the volume is primary in a remote mirror and copy relationship and the relationship is in a Full Duplex state, which means that the copy operation has completed

and the volume pair is synchronized, and that any updates to the primary volume are transferred synchronously to the secondary volume.

Primary Suspended
Indicates that the volume is primary in a remote mirror and copy relationship and the relationship is suspended, which means that the primary is no longer transferring data to the secondary, and any changed data that is at the primary is tracked in an out-of-sync bitmap.

" - " Indicates that there are no active volumes for the session.

SecondaryStatus
Indicates the secondary remote copy and mirror status of the volume. One of the following values is displayed:

Secondary Simplex
Indicates that the volume is not part of a remote mirror and copy relationship.

Secondary Copy Pending
Indicates that the volume is secondary in a remote mirror and copy relationship and the relationship is in a Copy Pending state, which means that the source and target volume are out-of-sync. In this situation, data still needs to be copied from the source to the target volume.

Secondary Full Duplex
Indicates that the volume is secondary in a remote mirror and copy relationship and the relationship is in a Full Duplex state, which means that the copy operation has completed and the volume pair is synchronized, and that any updates to the secondary volume are transferred synchronously from the primary volume.

Secondary Suspended
Indicates that the volume is secondary in a remote mirror and copy relationship and the relationship is suspended, which means that the primary is no longer transferring data to the secondary, and any changed data that is at the primary is tracked in an Out of Sync bitmap.

" - " Indicates that there are no active volumes for the session.

FirstPassComplete
Indicates whether the first cycle of the volume in the global mirror relationship has ended. The value displayed is either True or False.

AllowCascading
Indicates whether the volume can be a secondary in a remote mirror and copy relationship. The value displayed is either Enabled or Disabled.

SecondaryVolume
Indicates the volume ID that is represented as four hexadecimal digits in the form of *XYZZ*, where:

X Specifies the address group, 0 - F.

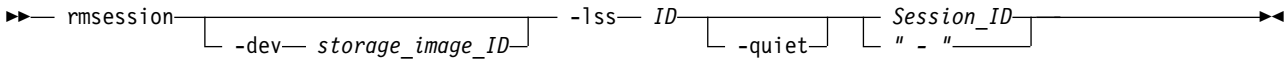
XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

The volume ID is always displayed as a full volume ID. However, if there are no active volumes in the Global Mirror session or if secondary information is not available, the output displays null (-).

rmsession

The **rmsession** command closes an existing Global Mirror session.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-lss *ID*
(Required) Specifies the logical subsystem (LSS) ID for the Global Mirror session that is being closed. The format of the LSS ID is a hexadecimal value in the range 00 - FE.

This parameter accepts a fully qualified LSS ID, which includes the storage image ID, or shortened version without the storage image ID if the **-dev** parameter is specified.

-quiet
(Optional) Turns off the end Global Mirror session confirmation prompt for this command.

Session_ID | **-**
(Required) Specifies the session ID on which Global Mirror processing is to be closed. A session ID is a hexadecimal number in the range 01 - FF.

If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive mode.

Example of a session ID: 01

Example

Invoking the rmsession command

```
dscli> rmsession -dev IBM.2107-75FA120 -lss 10 01
```

The resulting output

```
Are you sure you want to close Session ID 01? y/n Y
Global Mirror Session ID 01 closed successfully.
```

Framework commands

User interface framework commands are provided for the DS command-line interface.

The following framework commands are available:

dscli Starts the DS CLI. Use this command to complete storage management tasks from the command line.

echo Specifies whether the dscli echoes each specified command, or to display a user-specified string. If echo is turned on, each command is printed before it is processed. The echo can also be turned on or off using the dscli profile file or the **setenv** command.

exit Ends an interactive command-line interface session.

help Displays a list of commands available in a command-line interface and optionally displays the

syntax or brief description of each command. If you specify this command with no parameters, this command displays only a list of available commands.

helpmsg
Used to obtain details about information, warning, and error messages.

quit Ends an interactive command-line interface session.

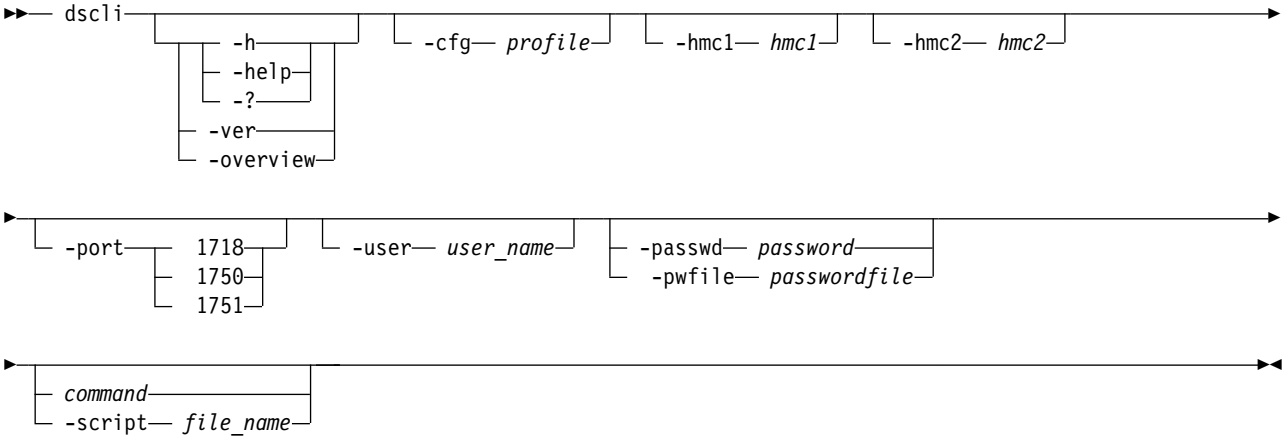
setenv Sets DS CLI environment variables that are located in the DS CLI profile.

showenv
Displays the DS CLI environment variables.

ver Displays the versions of the command-line interface, storage management console (HMC or SMC), and licensed machine code. It does not display the version number of the graphical user interface (GUI).

dscli

The **dscli** command starts the DS CLI. Use this command to run DS CLI commands in interactive, single-shot, or script mode.



Parameters

-help | -h | -?
(Optional) Displays a help screen about how to use the DS CLI program.

-ver
(Optional) Displays the DS CLI version.

-overview
(Optional) Provides overview information about the DS CLI.

-cfg *profile*
(Optional) Specifies a profile file. This parameter is not required if you use default profiles. The default profile file name is dscli.profile, and it is provided as part of the DS CLI package under the profile directory. See "Creating a default CLI profile" on page 10 for more information.

-hmc1 *hmc1*
(Optional) Specifies the primary management console IP address or the DNS name.

Note: You can connect the DS CLI to the HMC by using an IPV6 IP address if your Java virtual machine level supports IPV6.
This parameter is not required if you created this information as a profile variable.

-hmc2 *hmc2*
(Optional) Specifies the secondary management console IP address or the DNS name. This parameter is not required if you created this information as a profile variable.

-port 1718 | 1750 | 1751
(Optional) Specifies that the DS CLI should use only the specified port. By default, the DS CLI first attempts to connect by using port 1751 with an NIST-compliant certificate. If that connection fails, the DS CLI attempts to connect to an existing DS8000 port 1750 with the legacy certificate. However, checking multiple ports can cause a connection delay when a DS CLI connects to a DS8000 system that does not listen on the 1751 port. If the target DS8000 system is known, you can use this parameter to eliminate any connection delay that is caused by checking multiple ports.

1751 Connect by using port 1751 only on (DS8000 with NIST SP 800-131a-compliant certificate).

-user *user_name*
(Optional) Specifies your user name for entering DS CLI commands.

-passwd *password*
(Optional and not recommended) Specifies the password to authenticate when you start the CLI session. This parameter is not required nor recommended. If you use this method to designate your password, the password is displayed on the screen. Another option is to specify a password file (see the next parameter) that is used when you start the DS CLI.

-pwfile *passwordfile*
(Optional) Specifies a password file containing your password as an alternative to the **–passwd** parameter.

Note: You cannot specify both the **-pwfile** and the **-passwd** parameter when you start the DS CLI, otherwise an error message is given.

command
(Optional) Specifies the single command that you want to run.

Note: You cannot specify both the *command* and **-script** *file_name* parameters. Use **-script** for script mode, *command* for single-shot mode, and use neither for interactive mode.

-script *file_name*
(Optional) Initiates the script mode so that multiple dscli program commands are issued consecutively by using a saved file. Format options that are specified using the framework **setoutput** command apply to all commands in the file. Output from successful commands routes to stdout, and output from failed commands routes to stderr. If an error occurs during the processing of one of the commands in the file, the script exits at the point of failure and returns to the system prompt.

echo

The **echo** command allows you to specify whether the dscli will echo each specified command, or to display a user-specified string.

If echo is turned on, each command will be printed before it is executed. The echo can also be turned on or off using the dscli profile file or the **setenv** command.



Parameters

on | off | message
(Optional) Specifies whether the echo is turned on or off, or if a message is displayed. If it is turned

on, each command will be displayed before it is executed. If it is turned off, the dscli will not display anything before executing a command. You can also choose to echo a string of your choice. For example, if you enter `echo something else`, the string "something else" will be displayed. Issuing the **echo** command without any parameters will display whether **echo** is on or off.

Note: If the **echoprefix** variable is set, its value will be printed on the same line preceding the commands that are echoed.

Example

Issuing the echo command:

```
dscli> echo on
```

The resulting output

```
dscli> echo on
```

CMUC00212I echo: completed successfully.

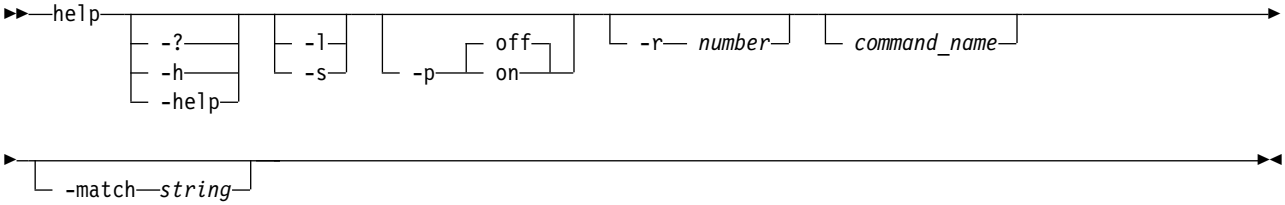
exit

The **exit** command ends an interactive command-line interface session.



help

The **help** command displays a list of commands available in a command-line interface and optionally displays the syntax or brief description of each command. If you specify this command with no parameters, this command displays only a list of available commands.



Parameters

-? | -h | -help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

-l
Displays a list of available commands with the syntax diagrams for each. If you specify a command name with this parameter, this command displays the syntax for only the specified command.

-s
Displays a list of available commands with a brief description of each. If you specify a command name with this parameter, this command displays a brief description for only the specified command.

-p *off | on*
Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This value is the default.

on Displays one page of text at a time. Pressing any key displays the next page.

-r *number*
Specifies the number of rows per page to display when the -p parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

command_name
Displays help information for the specified command, including the syntax diagram, parameter descriptions, return codes and errors, descriptions, examples, and miscellaneous remarks.

-match *string*
Displays a list of commands that contain the specified string. This parameter cannot be used with any other parameters.

Example

Invoke help

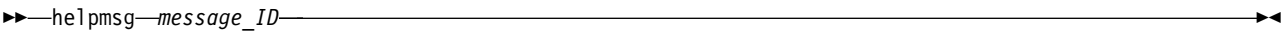
```
#dscli> help -s exit
```

The resulting output

```
Ends a command-line interface session.
```

helpmsg

The **helpmsg** command is used to obtain details about information, warning, and error messages.



Parameters

Notes:

1. The message information revealed by this command is a snapshot of what is available in your current code version.
2. This command does not work for GUI messages.
3. For the most up-to-date information, see the list of individual messages in the DS8000 online documentation.

message_ID
(Required) Specifies the message number that you are querying. You must enter the entire message number (for example, CMUC00246I, CMUC00244W, CMUC00247E). You do not have to enter all caps. Substitute keys are not allowed and cause an error if used.

Example

Invoking the helpmsg command

```
dscli> helpmsg cmuc00247e
```

The resulting output

```
CMUC00247E COMMAND_INFO You are attempting to create volumes for a
logical subsystem that does not match the rank group of the extent
pool that you specified.

Explanation
You create a logical volume from an extent pool in a rank group.
The logical volume is also a member of a logical subsystem group.
The logical subsystem group identifier and the rank group
identifier of the extent pool that you specify must be identical.
Even-numbered logical subsystem identifiers are associated with
rank group identifier 0, and odd-numbered logical subsystem
```

```
identifiers are associated with rank group identifier 1.

Action
Retry the create volume task and specify an extent pool and a
logical subsystem group that are in the same rank group.
```

quit

The **quit** command ends an interactive command-line interface session.

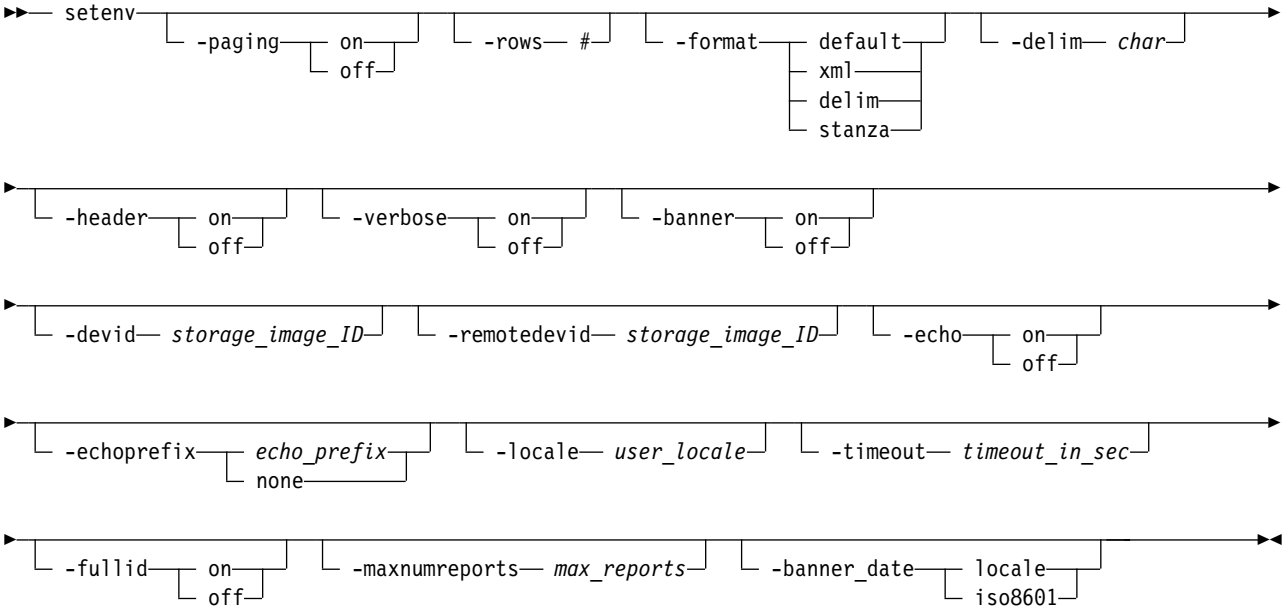


Parameters

-? | -h | -help
Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

setenv

The **setenv** command overrides the default values of the DS CLI environment variables. Default values are stored in the DS CLI profile. New values are valid for the current interactive session only and not permanently saved in the DS CLI profile.



Parameters

-paging on | off
(Optional) Specifies whether to display one page of text at a time or all of the text at once.

off
Displays all of the text at one time. This value is the default.

on Displays one page of text at a time. Pressing any key displays the next page.

-rows *#*
(Optional) Specifies the number of rows per page to display when **-paging** is set to *on*. The default is 24 rows. You can specify a value of 1 - 100.

-format default | **xml** | **delim** | **stanza**
(Optional) Specifies the format of the output. You can specify one of the following values:

default
Specifies that the output is to be displayed in a tabular format with a space as the delimiter between the columns. This value is the default.

delim
Specifies that the output format is to be set to a table and sets the column delimiter to a single character that is specified by the **-delim** *char* parameter.

xml
Specifies that the output is to be displayed in XML format.

stanza
Specifies that the output is to be displayed in a stanza (vertical table) format.

-delim *char*
(Optional) Specifies the delimiter character (such as a comma) used in the report.

-header on | **off**
(Optional) Specifies whether to display the table header.

on Displays the table header. This value is the default.

off
Does not display the table header.

-verbose on | **off**
(Optional) Specifies whether to enable verbose mode.

off
Disables verbose mode. The value that you specify takes the place of the value that you specified in the profile.

on Enables verbose mode.

-banner on | **off**
(Optional) Specifies whether the banner (command header) message is enabled.

off
Turns off the header mode so that the command header does not display.

on Turns on the header mode so that the command header is displayed.

-devid *storage_image_ID*
(Optional) Specifies the default devid.

-remotedevid *storage_image_ID*
(Optional) Specifies the default remote devid.

-echo on | **off**
(Optional) Specifies whether the echo is turned on or off. If it is turned on, each command is printed before it is run.

-echoprefix *echo_prefix* | **none**
(Optional) Specifies a prefix to display before each echoed command. `-echoprefix none` indicates that no prefix is displayed.

-locale *user_locale*
(Optional) Specifies the output language on the local computer. The default locale is based on the user environment settings. Example: en

-timeout *timeout_in_sec*
(Optional) Specifies the timeout value of client/server synchronous communication. The timeout value is displayed in seconds.

-fullid on | off
(Optional) Specifies that IDs are displayed in the fully qualified format, which includes the storage image ID. The default value is *off*.

-maxnumreports *max_reports*
(Optional) Specifies the maximum number of records for the performance report. The default value is 256.

-banner_date locale | iso8601
(Optional) Specifies the banner date and time format as follows:

locale
Specifies the locale format as set with the **-local** parameter, or by default, from the OS locale setting.

iso8601
Specifies the ISO 8601 date/time format as follows:

```
yyyy-MM-dd'T'HH:mm:ssZ
```

where:

yyyy the year

MM the month (01-12)

dd the day (01-31)

HH the hour (00-23)

mm the minutes (00-59)

ss the seconds (00-59)

Z the time zone offset from UTC [-HHmm | +HHmm]

The default value is *locale*.

Example

Enter the **setenv** command:

Invoking the setenv command

```
dscli> setenv -echo on
```

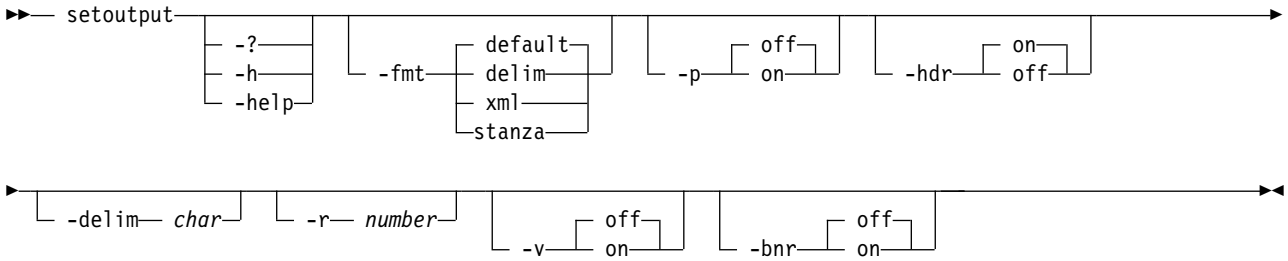
The resulting output

CMUC00043I setenv: Environment variable successfully set.

setoutput

The **setoutput** command sets or displays command output format options.

Use this command to set either default or user-defined output formats. The output format remains for the during the active command session unless reset either with a command option or the resubmission of the **setoutput** command. Running this command with no parameters displays the current output settings in the default output format.



Parameters

-? | -h | -help
(Optional) Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

-fmt default | delim | xml | stanza
(Optional) Specifies the format of the output. You can specify one of the following values:

default
Specifies that the output is to be displayed in a tabular format with a space as the delimiter between the columns. This value is the default.

delim
Specifies that the output format is to be set to a table and sets the column delimiter to a single character that is specified by the **-delim** *char* parameter.

xml
Specifies that the output is to be displayed in XML format.

stanza
Specifies that the output is to be displayed in a stanza (vertical table) format.

-p off | on
(Optional) Specifies whether to display one page of text at a time or all of the text at once.

off
Displays all of the text at one time. This value is the default.

on Displays one page of text at a time. Pressing any key displays the next page.

-hdr on | off
(Optional) Specifies whether to display the table header.

on Displays the table header. This value is the default.

off
Does not display the table header.

-delim *char*
(Optional) Specifies that the delimiter character (such as a comma) used in the report.

-r *number*
(Optional) Specifies the number of rows per page to display when the **-p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

-v off | on
(Optional) Specifies whether to enable verbose mode.

off
Disables verbose mode. This value is the default.

on Enables verbose mode.

-bnr off | on

(Optional) Specifies whether the banner (command header) message is enabled.

off

Turns off the header mode so that the command header does not display.

on Turns on the header mode so that the command header is displayed.

Format Examples

Invoke the setoutput command with no options

The **setoutput** command with no options displays the current output settings in the default format (space-separated plain text table), regardless of the values of the output settings.

```
dscli> setoutput
```

The resulting output

```
Paging Rows Format  Headers Verbose Banner
=========================================
Off    -    Default On      Off     On
```

Invoke the setoutput command using the -delim parameter

The following lines are an example of the commands that you can enter to get (long) output in comma-separated format for an unassigned rank only. Enter the **setoutput** command to specify the report format and then enter the **lsrank** command to designate the rank you want to query.

```
dscli> setoutput -fmt delim –delim ,
dscli> lsrank –dev IBM.2107-75FA120 -state unassigned
```

The resulting output

Note: Although this example shows the header turned on, you can choose to turn off the header. To turn off the header, issue the command and include the **-hdr** *off* parameter.

```
ID,Group,State,datastate,Array,RAIDtype,extpoolID,stgtype
=========================================================
R0,-,Unassigned,Normal,A0,5,-,fb
```

Invoke the setoutput command using the -fmt xml parameter

The following lines are an example of the commands that you can enter to get (long) output in XML format for an unassigned rank only. Enter the **setoutput** command to specify the report format and then enter the **lsrank** command to designate the unassigned rank that you want to query.

```
dscli> setoutput -fmt xml
dscli> lsrank –dev IBM.2107-75FA120 –state unassigned
```

The resulting output

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="CliRankHandler"><PROPERTY NAME="rank_id">
<DISPLAY TYPE="string">R0</DISPLAY><VALUE TYPE="string">R0
</VALUE></PROPERTY><PROPERTY NAME="grp"><DISPLAY TYPE="unit8">-
</DISPLAY><VALUE TYPE="unit16">-</VALUE></PROPERTY>
<PROPERTY NAME="state"><DISPLAY TYPE="string">Unassigned</DISPLAY>
<VALUE TYPE="string">unassigned</VALUE></PROPERTY>
<PROPERTY NAME="data"><DISPLAY TYPE="string">Normal</DISPLAY>
<VALUE TYPE="string">Normal</VALUE></PROPERTY>
<PROPERTY NAME="array_id><DISPLAY TYPE=string">A0
</DISPLAY><VALUE TYPE="string">A0</VALUE></PROPERTY>
<PROPERTY NAME="raidtype"><DISPLAY TYPE="unit8">5</DISPLAY>
<VALUE TYPE="string">5</VALUE><PROPERTY>
```

```
<PROPERTY NAME="extpool_id"><DISPLAY TYPE="string">-
</DISPLAY><VALUE TYPE=string">-<VALUE><PROPERTY>
<PROPERTY NAME="stgtype"><DISPLAY TYPE="string">fb</DISPLAY>
<VALUE TYPE="string">fb</VALUE><PROPERTY><INSTANCE>
<IRETURNVALUE>
```

Invoke the setoutput command using the -fmt stanza parameter

When columns are horizontally long, output can be difficult to visually align. The stanza format option eliminates this problem. The following output is an example of the commands that you can enter to get (long) output in stanza format for an unassigned rank only. Issue the **setoutput** command to specify the report format and then enter the **lsrank** command to designate the unassigned rank that you want to query.

```
dscli> setoutput -fmt stanza
dscli> lsrank –dev IBM.2107-75FA120 –state unassigned
```

The resulting output

```
ID        R0
Group     -
State     unassigned
datastate normal
Array     A0
RAIDtype  5
extpoolID -
stgtype   fb
```

showenv

The **showenv** command displays the DS CLI environment variables.

Example

Enter the **showenv** command.

```
dscli> showenv
```

The resulting output

```
paging        off
rows          24
format        default
delim         ,
header        on
verbose       off
banner        on
devid         IBM.2107-1301441
remotedevid   IBM.2107-1301441
echo          off
echoprefix    -
locale        zh
timeout       -
fullid        off
bannerDate    locale
maxNumReports 256
```

Report field definitions

paging

Indicates whether paging is turned on or off.

rows Indicates the number of rows per page that are displayed.

format
Indicates the output format.

delim Indicates the delimiter character that is used in the *delim* output format.

header
Indicates whether the header is turned on or off.

verbose
Indicates whether verbose mode is turned on or off.

banner
Indicates whether the banner is turned on or off.

devid Indicates the default devid.

remotedevid
Indicates the default remote devid.

echo Indicates whether the echo is turned on or off. If it is turned on, each command is printed before it is run.

echoprefix
Indicates that a prefix is displayed before each echoed command. A dash "-" is shown and no prefix is displayed if echoprefix is not specified.

locale Indicates the output language on the local computer.

timeout
Indicates the timeout value of client/server synchronous communication. The timeout value is displayed in seconds.

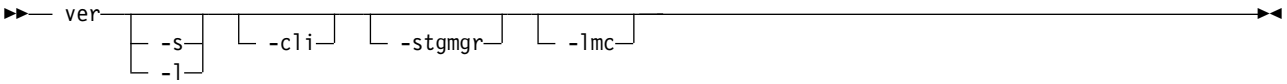
fullid Indicates that IDs are displayed in the fully qualified format, which includes the storage image ID.

bannerDate
Indicates the format of the banner date and time.

maxnumreports
Indicates the maximum number of records for the performance report.

ver

The **ver** command displays the versions of the command-line interface, storage management console (HMC or SMC), and licensed machine code. It does not display the version number of the Graphical User Interface (GUI).



Parameters

-s
(Optional) Specifies that the system displays the version of the command-line interface program. You cannot use the **-s** and **-l** parameters together.

-l
(Optional) Specifies that the system displays the following version information:

- The version of the DS command-line interface
- The version of the Storage Management GUI (SMC or HMC)

- The version of the licensed machine code
- The storage image ID, which consists of the manufacturer, machine type, and serial number (MTS)
- The version of the DS CLI that is installed on the HMC.
- The version of the currently active code bundle that is installed.

You cannot use the **-s** and **-l** parameters together.

-cli
(Optional) Displays the version of the DS Command-line interface program. Version numbers are in the format *version.release.modification.fix level*.

-stgmgr
(Optional) Displays the HMC code version number. This value is not the version number of the Graphical User Interface (GUI).

-lmc
(Optional) Displays the version of the DS licensed machine code.

Example

Invoking the ver command

```
dscli> ver –l
```

The resulting output

```
DS Date/Time: February 27, 2014 6:00:01 AM MST IBM DSCLI
 Version: 7.7.30.198 DS:

DSCLI           7.7.30.198
StorageManager  7.7.7.0.20130628.1
HMC DSCLI       7.7.30.180
====================Version===================
Storage Image      LMC          Bundle Version
==============================================
IBM.2107-75YZ881   7.7.30.128   87.30.40.0
```

Report field definitions

DSCLI*
Indicates the version of the DSCLI, in this format: *version.release.modification.fix level*.

StorageManager+
Indicates the version of the Storage Manager (HMC microcode). This value is not the version number of the DS Storage Manager (Graphical User Interface, GUI).

HMC DSCLI+
Indicates the version of the HMC DSCLI in this format: *version.release.modification.fix level*.

Storage Image+
Indicates the storage image ID, which consists of the manufacturer, machine type, and serial number (MTS).

LMC+ Indicates the version of the Licensed Machine Code.

Bundle Version+
Indicates the version of the currently active installed code bundle.

Key:

***** Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

Install software commands

Use the install software commands to install software and display software information.

The following install software commands are available:

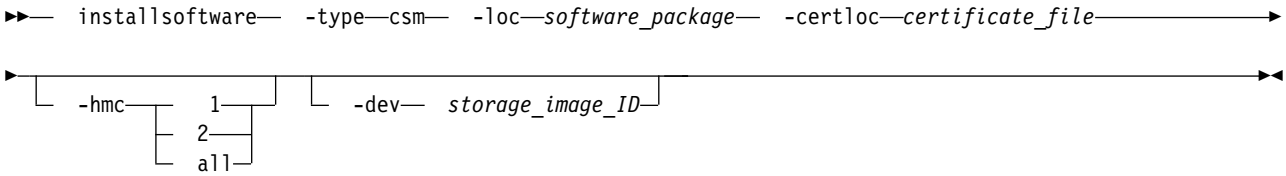
installsoftware
Installs software on the storage server.

lssoftware
Displays a list of installed software information.

installsoftware

The **installsoftware** command installs software for the storage server.

Note: You can specify the software that was previously downloaded from the IBM Support Portal. The Support Portal for CSM is available from https://www.ibm.com/support/entry/myportal/product/system_storage/storage_software/storage_infrastructure_management/ibm_copy_services_manager?productContext=332956944



Parameters

-type
(Required) Specifies the software type of the installation package to be installed. Currently, the only supported type is CSM (IBM Copy Services Manager).

-loc *software_package*
(Required) Specifies the full path of the specified software installation package to be installed. For example, `csm-setup-6.2.0-linux-x86_64.bin`.

Note: Ensure that no spaces are included in the path you specify for the location of the software package.

-certloc *certificate_location*
(Required) Specifies the full path of the certificate file location. For example, `C:\Downloads\CSM_Linux\csm-setup-6.2.0-linux-x86_64.bin -certloc`.

Note: Ensure that no spaces are included in the path you specify for the location of the certificate file.

-hmc 1 | 2 | all
(Optional) Specifies the primary HMC hmc 1 or secondary HMC hmc 2 where the software is installed. The default is all. For example, `C:\Downloads\CSM_Linux\csm-setup-6.2.0-linuxx86_64.bin.crt -hmc 1`.

Note: -hmc 1 specifies the primary HMC and -hmc 2 specifies the secondary regardless of how -hmc 1 and -hmc 2 were specified at the **dscli** startup process.

A **dscli** connection can succeed even if you inadvertently specify the primary HMC with -hmc 2 and the secondary backup HMC with -hmc 1 at the **dscli** command startup. In this case, if you specify -hmc 1 with this command, the software is downloaded to the actual primary HMC.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID. A storage image ID consists of a manufacturer, machine type, and serial number. For example, 2107-75FA120.

Example

Invoking the installsoftware command

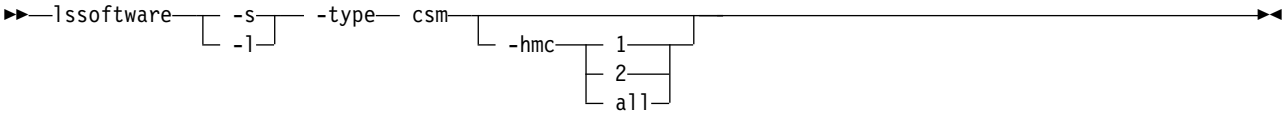
```
dscli> installsoftware -type csm -loc
 C:\Downloads\CSM_Linux\csm-setup-6.2.0-linux-x86_64.bin -certloc
 C:\Downloads\CSM_Linux\csm-setup-6.2.0-linux-x86_64.bin.crt -hmc 1
```

Output

```
CMUC00294I installsoftware: Upload file successfully.
CMUC00294I installsoftware: Software CSM is successfully installed on HMC 1.
```

lssoftware

The **lssoftware** command displays a list of installed software.



Parameters

-type
(Required) Specifies IBM Copy Services Manager is installed on this storage server.

-hmc 1 | 2 | all
(Optional) Specifies the primary HMC hmc 1 or secondary HMC hmc 2 where the software is downloaded.

Note: -hmc 1 specifies the primary HMC and -hmc 2 specifies the secondary HMC regardless of how -hmc 1 and -hmc 2 were specified at the **dscli** startup process.

For example, a **dscli** connection can succeed even if you inadvertently specify the primary HMC using -hmc 2 and the secondary backup HMC using -hmc 1 at the **dscli** startup process. In this case, if you specify -hmc 1 with this command, the software is downloaded to the actual primary HMC.

Example

```
dscli> lssoftware -l
Sun Aug 11 02:23:49 PST 2015 IBM DS CLI
Type         Version        Status        HMC
==========================================
CSM            6.1          Running        1
```

Report field definitions

Type
Indicates the software type.

Version
Indicates the currently installed software version.

Status
Indicates the status for the installed software.

Running
The software is enabled and is working normally.

Not running
The software is not enabled or is working abnormally.

HMC
Indicates the number of the HMC on which the software is installed.

1 Indicates the primary HMC.

2 Indicates the secondary HMC.

all
Indicates both primary and secondary HMCs.

I/O port and host connection configuration commands

Specific DS CLI commands are used to configure and display storage image I/O port information.

Storage image I/O port commands

There are specific commands used to configure and display storage image I/O port information.

The following storage image I/O port commands are available:

lsioport
Displays a list of I/O ports on a specified storage image and optionally provides performance metrics for each I/O port that is listed.

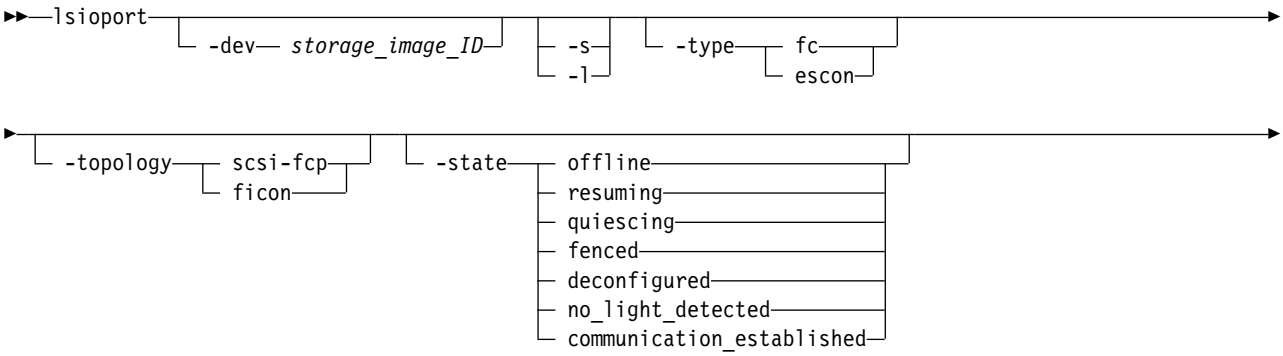
lsportprof
Displays a list of port profiles that are supported on a storage system and their recommended address discovery and logical block size values.

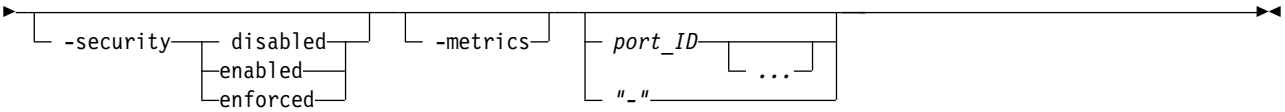
setioport
Configures one or more I/O ports for open systems or IBM Z host system connections.

showioport
Displays the properties of a specified I/O port. It optionally displays the performance metrics for the I/O port.

lsioport

The **lsioport** command displays a list of all I/O ports (ESCON and Fibre Channel type) that are installed in a specified storage image, and optionally provides performance metrics for each I/O port listed.





Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified port ID. It is also required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-s
(Optional) Displays fully qualified port IDs. You cannot use the **-l** and **-s** parameters together.

-l
(Optional) Displays default output plus the I/O port interface speed.

-type fc | escon
(Optional) Displays I/O ports of the specified port type. The *Escon* option can be specified as the port type

-topology scsi-fcp | ficon
(Optional) Displays Fibre Channel I/O ports with the specified topology.

-state offline | resuming | quiescing | fenced | deconfigured | no_light_ detected | communication_established
(Optional) Displays I/O ports of the specified state. See the Example section for the state descriptions.

-security enabled | disabled | enforced
(Optional) Displays the current security setting of an I/O port. One of the following values is displayed.

Disabled
Endpoint Security is not enabled for the port.

Enabled
Endpoint Security is enabled to allow authentication between the host and the port, but authentication is not required for I/O processing.

Enforced
Endpoint Security is enabled and authentication is required between the host and the port for I/O processing. Host ports that do not support security or do not complete authentication will be logged out and denied login.

-metrics
(Optional) Displays port ID and performance metrics for each port that is specified.

Note: All performance counts are an accumulation since the most recent counter-wrap or counter-reset operation. I/O port performance counters are reset with a storage system power-on sequence.

port_ID ... | **-**
(Optional) Displays I/O ports that match the specified IDs. This parameter accepts a fully qualified port ID, which includes the storage image ID. You can also specify a shortened version without the storage image ID when the **-dev** parameter is specified.

A port ID is prefixed with the letter "I" and consists of four hexadecimal characters in the format *EEAP*, where:

- *EE* is an I/O port enclosure number in the range of 00 - 17.
- *A* is the adapter number and is specified as 1, 2, 4, or 5.
- *P* is the port number (0 - 3).

Example 1

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

```
dscli> lsioport –dev 2107-75FA120 -l
```

ID	WWPN	State	Type	Topo	Portgrp	Speed	Security
2107-75FA120/I0111	307BCF30 A3299E0	Communication Established	Fibre Channel-LW	SCSI-FCP	0	1 Gb/s	Disabled

Frame	I/O Enclosure	HA Card
1	2	4

Report field definitions (without the -metrics parameter)

ID Indicates the fully qualified port ID, which includes the storage image ID.

WWPN
Indicates the Fibre Channel worldwide port number. If the port type is not Fibre Channel, the value that is displayed is " - ".

State
Indicates the current I/O port status. One of the following values can be displayed:

Offline
Indicates the state of the port that is internal to the storage system and indicates that the port is not capable of processing I/O.

Resuming
Indicates that the storage system is coming online.

Quiescing
Indicates that the storage system is going offline.

Fenced
Indicates that the storage system failed and is offline.

Deconfigured
Indicates that the I/O port is being deleted.

No Light Detected
Indicates that there is no light that is detected on the port and, therefore, the port is not capable of processing I/O.

Communication Established
Indicates that the storage system can process all functions.

Type
Indicates the port type. The following values can be displayed:

- Fibre Channel-SW (short wave)
- Fibre Channel-LW - (long wave, 10 KM)
- Fibre Channel-LW 4 KM
- ESCON

Topo
Indicates the I/O port topology. The following values can be displayed:
- SCSI-FCP
- FICON
- " - " This value is displayed when the port type is not Fibre Channel.

Portgrp
Indicates the identifier that associates a subset of the I/O ports that are operating in anonymous access mode. Default value is 0 when these subsets are not specified.

Speed
Indicates the I/O port interface speed. The following values can be displayed:
- ESCON ports = 200 Mb/s
- FCP ports = 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s, 16 Gb/s, 32 Gb/s
- FICON ports = 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s, 16 Gb/s, 32 Gb/s

Security
Indicates the current security setting of an I/O port. One of the following values is displayed.

Disabled
Endpoint Security is not enabled for the port.

Enabled
Endpoint Security is enabled to allow authentication between the host and the port, but authentication is not required for I/O processing.

Enforced
Endpoint Security is enabled and authentication is required between the host and the port for I/O processing. Host ports that do not support security or do not complete authentication will be logged out and denied login.

Frame
Identifies the unique identifier of the storage system frame that contains the I/O enclosure.

I/O Enclosure
Identifies the unique identifier of the I/O enclosure that contains the host adapters.

HA Card
Identifies the ID number of the host adapter.

Report field definitions (with the -metrics parameter)

When you use the **-metrics** parameter and do not specify a port type, separate reports are displayed for DS8000 models:
- FICON/ESCON IO port type: A banner is displayed (for example: ===FICON/ESCON-Enabled I/O Ports===) before each report.
- SCSI-FCP IO port type.
- Security Login Counts per FC IO Port

Note: A report is not displayed for a port type that has no enabled ports.

FICON/ESCON ports

Each of the following headers and value types is displayed:

ID Indicates the fully qualified port ID.

Date
Indicates the current time stamp for the I/O port performance counters. For example, 08/11/14 02:23:49 is the format that is used to report this value.

bytewrit
Indicates the number of bytes that are written in 128 KB increments.

byteread
Indicates the number of bytes that are read in 128 KB increments.

Reads
Indicates a value that is based on extended count-key-data (ECKD™) data received operations.

Writes
Indicates a value that is based on ECKD data transferred operations.

Timewrite
Indicates a value that is based on the ECKD data that is transferred (write-accumulated time) on a channel. The displayed value is based on increments of 16 milliseconds.

Timeread
Indicates a value that is based on the ECKD data that is received (read-accumulated time) on a channel. The displayed value is based on increments of 16 milliseconds.

CmdRetries
Indicates the number of retries that are requested for transport mode write operations.

TransferReady
Indicates the number of transport mode operations from channels that require transfer ready and that were received from channels that support transport mode command retry operations.

SCSI-FCP ports

Each of the following headers and value types is displayed:

ID Indicates the fully qualified port ID.

Date
Indicates the current time stamp for the I/O port performance counters. For example, 08/11/05 02:23:49 is the format that is used to report this value.

Bytewrit
Indicates a value for the remote mirror and copy data transferred operations in increments of 128 KB.

Byteread
Indicates a value for the remote mirror and copy data received operations in increments of 128 KB.

Writes
Indicates a value for the remote mirror and copy data transferred operations.

Reads
Indicates a value for the remote mirror and copy data received operations.

Timewrite
Indicates a value that is based on the remote mirror and copy data transferred (write-accumulated) time on a channel. The displayed value is based on increments of 16 milliseconds.

Timeread
Indicates a value for the remote mirror and copy data received (read-accumulated) time on a channel. The displayed value is based on increments of 16 milliseconds.

Byteread
Indicates a value that is based on the SCSI data received operations. The displayed value is based on increments of 128 KB.

Reads
Indicates a value that is based on the SCSI data transferred operations.

Writes
Indicates a value that is based on the SCSI data transferred operations.

Timeread
Indicates a value that is based on the SCSI data received (read-accumulated) time on a channel. The displayed value is based on increments of 16 milliseconds.

Timewrite
Indicates a value that is based on the SCSI data transferred (write-accumulated) time on a channel. The displayed value is based on increments of 16 milliseconds.

Security Login Counts per FC IO Port

Note: The Security Login Counter per FC IO Port report applies to all types of ports.

Each of the following headers and value types is displayed.

ID The fully qualified port ID.

Logins
The number of remote hosts or peer *N*_ports currently logged in to the specified I/O port.

SecCapableLogins
The number of logins that are capable of Fibre Channel security protocols. (The security bit is set in both *N*_ port login and response.) If the port is security-enabled, the remote ports are expected to perform authentication.

AuthLogins
The number of security-capable logins that have negotiated link authentication.

EncryptedLogins
The number of remote host or peer *N*_ports currently logged into the specified port that successfully authenticated and negotiated an encryption algorithm for IBM Fibre Channel Endpoint Security.

Example 1

```
dscli> lsioport
```

ID	WWPN	State	Type	Topo	Portgrp	Security
I0100	50050763 06081343	Communication Established	Fibre Channel-SW	FICON	0	Disabled

Example 2

```
dscli> lsioport -security enabled
```

ID	WWPN	State	Type	Topo	Portgrp	Security
I0101	50050763 06085344	Communication Established	Fibre Channel-SW	SCSI-FCP	0	Enabled

Example 3

```
dscli> lsioport -metrics I0100
```

ID	Logins	SecCapableLogins	AuthLogins	EncryptedLogins
I0101	0	0	0	0

Report field definitions (with -metrics and –l parameters)

Fibre Channel Link Errors: Each of the following headers and value types is displayed:

ID Indicates the fully qualified port ID.

LinkFailErr
Indicates the total number of miscellaneous Fibre Channel link errors.

LossSyncErr
Indicates the number of loss of synchronization errors. These errors occur when there is a confirmed and a persistent synchronization loss on the Fibre Channel link.

LossSigErr
Indicates the number of times that a loss of signal was detected on the Fibre Channel link when a signal was previously detected.

PrimSeqErr
Indicates the number of primitive sequence protocol error counts where an unexpected primitive sequence was received.

InvTxWordErr
Indicates the number of times a "bit" error was detected. Examples of a "bit" errors are a code violation, an invalid special code alignment, or a disparity error.

CRCErr
Indicates the number of times the CRC of a received frame is in error.

LRSent
Indicates the number of times the port changed from an active (AC) state to a Link Recovery (LR1) state.

LRRec
This count is the number of times the port that is changed from an active (AC) state to a Link Recovery (LR2) state.

IllegalFrame
Indicates the number of frames that violated the Fibre Channel protocol. One example of a violation is an invalid frame header, which occurs when the first frame of a data sequence is missing and a subsequent data frame is detected as illegal.

OutOrdData
Indicates the number of times that an out-of-order frame is detected. The frame is either missing from a data sequence or it is received beyond the sequence reassembly threshold of the port.

OutOrdACK
Indicates the number of times that an out-of-order ACK (ACKnowledgment field of the TCP protocol) frame is detected. The frame is either missing from a data sequence or it is received beyond the sequence reassembly threshold of the port.

DupFrame
Indicates the number of times a frame was received that was detected as previously processed.

InvRelOffset
Indicates the number of times that a frame was received with bad relative offset in the frame header.

SeqTimeout
Indicates the number of times the port detected a timeout condition after receiving a sequence initiative for a Fibre Channel exchange.

BitErrRate
Indicates the number of the bit error (invalid transmission word) bursts for the previous 5-minute counting window. This number is reset every 5 minutes.

Read Diagnostic Parameters: Each of the following headers and value types is displayed:

ID Indicates the fully qualified port ID.

TxPower
Indicates the measured coupled TX output power in dBm and uW. The maximum value is 18.1 dBm (6500 uW).

If value displays 0 when the link state is active, the function is not supported by the current level of microcode.

RxPower
Indicates the measured received optical power in dBm and uW. The maximum value is 18.1 dBm (6500 uW).

If value displays 0 when the link state is active, the function is not supported by the current level of microcode.

TransceiverTemp
Indicates the internally measured transceiver temperature. This value is reported in the range -128 C to + 128 C.

If value displays 0 when the link state is active, the function is not supported by the current level of microcode.

SupplyVolt
Indicates the internally measured supply voltage. This value is reported in mV and the value range is 0 - 6550 mV.

If value displays 0 when the link state is active, the function is not supported by the current level of microcode.

TxBiasCurrent
Indicates the measured transmitter laser bias current. This value is reported in mA and is in the value range of 0-131mA.

If value is 0 when the link state is active, the function is not supported by the current level of microcode.

ConnectorType
Indicates 1 of 3 connector type values:

SFP+
Specifies a transmit power level of small form factor pluggable device.

Unknown
Specifies that the connector type might not be determined.

Dash (-)
Specifies that the data was not available.

TxType
Indicates a value for the Port Tx type:

Laser-SW
Specifies a short-wave laser port.

Laser LC 1310-LW
Specifies a long-wave laser LC 1310 nm port.

Laser LL 1550-LW
Specifies a long-wave laser LotusLive™ 1550 nm port.

Non-Optical
Specifies that the port is not optical or is of another type.

Unknown
Specifies that the port type is not recognized.

Dash (-)
Specifies that the data was not available.

CurrentSpeed
Indicates the current operating link speed:

- Inactive – Link not operational
- 200 Mb/s – ESCON
- 1 Gb/s – Fibre Channel 1 Gb/s
- 2 Gb/s – Fibre Channel 2 Gb/s
- 4 Gb/s – Fibre Channel 4 Gb/s
- 8 Gb/s – Fibre Channel 8 Gb/s
- 16 Gb/s – Fibre Channel 16 Gb/s
- 32 Gb/s - Fibre Channel 32 Gb/s

Note: A dash (-) specifies that data was not available.

FECStatus
Indicates the Forward Error Correction (FEC) process status on the link:

Active
Specifies that FEC is active on the link.

Inactive
Specifies that FEC is not active on the link.

UncorrectedBlks(RDP)
Specifies the number of bad data blocks that were not corrected by the Forward Error Correction (FEC) process. For more information, see the **showioport** command.

Note: A dash (-) means that the data is not available.

CorrectedBlks(RDP)
Specifies the number of bad data blocks that were corrected by forward error correction (FEC) process. For more information, see the **showioport** command.

Note: A dash (-) means that the data is not available.

lsportprof

The **lsportprof** command displays a list of port profiles that are supported on a storage system and their recommended address discovery and logical block size values.

You can use this command to view known values for the block size (lbs) and address discovery (**addrdiscovery**) parameters in the **mkhostconnect** command.

Note: Use this command to get the recommended values for the **mkhostconnect** command.



Parameters

storage_image_ID | **-**
(Required) Displays a list of port profiles for the specified storage image IDs. A storage image ID consists of manufacturer, type, and serial number.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

For DS8000, example: IBM.2107-75FA120.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lsportprof** command.

Invoking the lsportprof command

```
dscli> lsportprof IBM.2107-75FA120
```

The resulting output

Name	AddrDiscovery	LBS
IBM pSeries - AIX	ReportLUN	512
IBM pSeries - pLinux	LUNPolling	512

Report column definitions

Name Specifies the name of the host connection behavior profile. The port profile specifies a host or operating system type.

AddrDiscovery
Specifies the address discovery method. One of the following values is displayed:

LUN Polling
Specifies that host system LUN access is limited to a maximum of 256 LUNs.

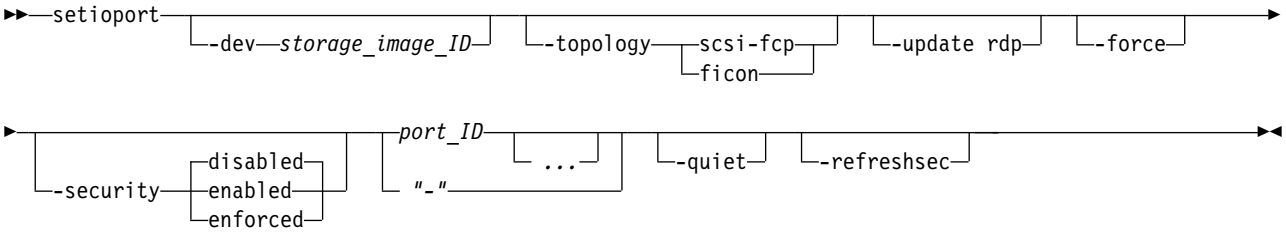
Report LUN
Specifies that host system LUN access is limited to a maximum of 64000 LUNs.

LBS Specifies the logical block size. One of the following values is displayed:

- 512 - This value is displayed for all hosts except OS400.
- 520 - This value is displayed for an OS400 host.

setioport

The **setioport** command configures one or more I/O ports for open systems or IBM Z host system connections. This command cannot be used for ESCON ports.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial

number. The storage image ID consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified port ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-topology scsi-fcp | ficon
(Optional) Specifies the topology for an I/O port, either SCSI-FCP or FICON.

scsi-fcp
The SCSI-FCP topology setting enables the SCSI ULP with a point-to-point or switched fabric topology. PPRC path I/O operations are enabled for this setting.

ficon
The Ficon topology setting enables the FICON ULP with a point-to-point or switched fabric topology. PPRC path I/O operations are not supported for FICON ULP.

-update rdp
(Optional) Forces the port to send a request to the attached port to retrieve the current Read Diagnostic Parameters (RDP) data from the port.

-force
(Optional) Allows a topology change to an online I/O port even if a current topology is set. Use this option if you want to change a topology after it is already set.

-security disabled | enabled | enforced
Specifies the security setting for an I/O port.

disabled
(Default) Security is not enabled for the specified I/O port.

enabled
Security is enabled for the specified I/O port. In this state, the port allows security protocols to negotiate authentication but will not enforce it's use before allowing I/O.

enforced
Security is enforced for the specified I/O port. In this state, the port requires secure login and the use of security protocols before allowing I/O. Ports that do not support security or do not complete authentication will be logged out and denied login.

port_ID ... | **-**
(Required) Specifies the I/O port ID. Accepts a fully qualified port ID, which includes the storage image ID, or a shortened version without the storage image ID when the **-dev** parameter is specified.

A port ID is prefixed with the letter *I* and consists of four hexadecimal characters in the format *EEAP*, where:

- *EE* is an I/O port enclosure number in the range of 00 - 17.
- *A* is the adapter number and is specified as 0, 1, 3, or 4.
- *P* is the port number (0 - 3).

To specify a range of port IDs, separate the port IDs with a hyphen.

You must separate multiple port IDs or ranges of port IDs by a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

-quiet
(Optional) Turns off the confirmation prompt for -security parameter.

-refreshsec
(Optional) Forces the host to re-authenticate with the specified I/O port.

Example 1

```
dscli> setioport -refreshsec I0301
CMUC00011I setioport: I/O Port I0301 successfully configured.
```

Example 2

```
dscli> setioport -security enabled I0230
CMUC00590W setioport: A temporary interruption of I/O will occur on this port for a few seconds
while its security is modified.
Are you sure that you want to modify the security of I/O Port I0230? [Y/N]: y
CMUC00011I setioport: I/O Port I0230 successfully configured.
```

Example 3

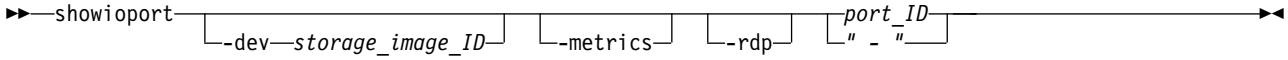
```
dscli> setioport -quiet -security disabled I0230
CMUC00011I setioport: I/O Port I0230 successfully configured.
```

Example 4

```
dscli> setioport -update rdp IBM.2107-I0201
CMUC00011I setioport: Read Diagnostic Parameters of I/O Port I0201 and its attached port
are successfully updated.
```

showioport

The **showioport** command displays properties of an I/O port. It optionally displays the performance metrics for a specific I/O port, or the Read Diagnostic Parameters (RDP) for a local port and an attached port.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified port ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-metrics
(Optional) Displays the port ID and the performance metrics for the specified I/O port.

Note: All performance counts are an accumulation since the most recent counter-wrap or counter-reset operation. I/O port performance counters are reset with a storage system power-on sequence.

-rdp
(Optional) Displays the Read Diagnostic Parameters (RDP) data for the specified local port and attached port. This data includes link statistics, current link speed and capabilities, buffer credit settings, and transceiver diagnostic information.

Note: This parameter is not available for I/O ports on 8Gb/s adapters. An error message is displayed to indicate if the function is not supported.

port_ID | **-**
(Required) Displays the property level details for the specified port IDs. This parameter accepts a fully qualified unique port ID that is represented in the following format: `manufacturer.machine type-serial number/portID`.

For example, for DS8000, IBM.2107-75FA120/I0110.

A port ID is prefixed with the letter *I* and consists of four hexadecimal characters in the format *EEAP,* where:

- *EE* is an I/O port enclosure number in the range of 00 - 17.
- *A* is the adapter number and is specified as 1, 2, 4, or 5.
- *P* is the port number (0 - 3).

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 1

Invoking the showioport command to show port properties

```
dscli> showioport -dev IBM.2107-75DMC81 I0000
ID            I0000
WWPN          500500763080005AE
State         Communication Established
Loc           U1400.1B1.1300025-P1-C1-T0
Type          Fibre Channel-SW
Speed         8 Gb/s
Topo          SCSI-FCP
Portgrp       0
UnkSCSIlog    -
Physloc       R1-I1-C1-T0
Security      Disabled
```

Report field definitions (with the -port parameter)

ID Specifies the fully qualified unique port ID.

WWPN
Specifies the Fibre Channel I/O port worldwide port number (WWPN). If the port type is not Fibre Channel, this value is specified as a " - ".

State
Specifies the current state of the I/O port. One of the following values is displayed:

Offline
Indicates the state of the port that is internal to the storage system and indicates that the port is not capable of processing I/O.

Resuming
Indicates that the storage system is coming online.

Quiescing
Indicates that the storage system is going offline.

Fenced
Indicates that the storage system failed and is offline.

Deconfigured
Indicates that the I/O port is being deleted.

No Light Detected
Indicates that there is no light detected on the port and, therefore, the port is not capable of processing I/O.

Communication Established
Indicates that the storage system can process all functions.

Loc
Specifies the storage enclosure location by identifying the storage unit frame that contains the storage enclosure. The location format is *Utttt.mmm.ppssss*.

Type
Specifies the port type. The following values can be displayed:
- Fibre Channel-SW - (SW stands for short wave)
- Fibre Channel-LW - (LW stands for long wave, 10 KM)
- Fibre Channel-LW 4 KM - (LW stands for long wave, 4 KM)
- ESCON

Speed
Specifies the I/O port interface speed. The following values can be displayed:
- ESCON ports = 200 Mb/s
- FCP ports = 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s , 16 Gb/s, 32 Gb/s
- FICON ports = 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s, 16 Gb/s, 32 Gb/s

Topo
Specifies the port topology. If the port type is not Fibre Channel, then the displayed value is " - ". One of the following values is displayed:
- SCSI-FCP
- FICON
- " - " (if not Fibre Channel)

Portgrp
Specifies an identifier that associates a subset of the I/O port objects that are operating in anonymous access mode.

UnkSCSIlog
Specifies a list of unknown SCSI N-port WWPN identifiers that attempted to log in to this I/O port.

Physloc
Specifies the physical location of the I/O port. The I/O port location code is a combination of the rack, I/O enclosure, card, and port that provides the physical link. The value of the location uses the following format: R(1-2)-I(1-8)-C(1-6)-P(1-4).
- R is the rack location
- I is the I/O enclosure
- C is the card
- P is the port of the adapter

Note: R1-I3-C2-P1 is an example of this format.

Security
Indicates the current security setting of an I/O port. One of the following values is displayed.

disabled
(Default) Security is not enabled for the specified I/O port.

enabled
Security is enabled for the specified I/O port. In this state, the port allows security protocols to negotiate authentication but will not enforce it's use before allowing I/O.

enforced
Security is enforced for the specified I/O port. In this state, the port requires secure login and the

use of security protocols before allowing I/O. Ports that do not support security or do not complete authentication will be logged out and denied login.

Example 2

Invoking the showioport command to show port performance metrics

```
dscli> showioport -dev IBM.2107-75DMC81 -metrics I0000
ID                          I0000
Date                        04/17/2017 11:02:45 MST
byteread (FICON/ESCON)      0
bytewrite (FICON/ESCON)     0
Reads (FICON/ESCON)         0
Writes (FICON/ESCON)        0
timeread (FICON/ESCON)      0
timewrite (FICON/ESCON)     0
bytewrit(PPRC)              0
byteread PPRC)              0
Writes(FICON/ESCON)         0
timeread (FICON/ESCON)      0
timewrite(FICON/ESCON)      0
CmdRetries (FICON)          0
TransferReady(FICON)        0
bytewrit (PPRC)             0
byteread (PPRC)             0
Writes (PPRC)               0
Reads (PPRC)                0
timewrite (PPRC)            0
timeread (PPRC)             0
byteread (SCSI)             1271326535
bytewrit (SCSI)             65721755
Reads (SCSI)                2479310687
Writes (SCSI)               1170136080
timeread (SCSI)             81426674
timewrite (SCSI)            407813158
LinkFailErr (FC)            8
LossSyncErr (FC)            45
LossSigErr (FC)             0
PrimSeqErr (FC)             0
InvTxWordErr (FC)           125
CRCErr (FC)                 0
LRSent (FC)                 0
LRRec (FC)                  0
IllegalFrame (FC)           0
OutOrdData (FC)             0
OutOrdACK (FC)              0
DupFrame (FC)               0
InvRelOffset (FC)           0
SeqTimeout (FC)             0
BitErrRate (FC)             0
RcvBufZero (FC)             0
SndBufZero (FC)             0
RetQFullBusy (FC)           0
ExchOverrun (FC)            0
ExchCntHigh (FC)            0
ExchRemAbort (FC)           0
CurrentSpeed (FC)           0
%UtilizeCPU (FC)            0
TxPower (RDP)               -2.8 dBm (525.2 uW)
RxPower (RDP)               -3.6 dBm (440.5 uW)
TransceiverTemp (RDP)       65 C
SupplyVolt (RDP)            3304.7 mV
TxBias (RDP)                8.484 mA
Connector Type (RDP)        SFP+
TxType (RDP)                Laser-SW
FEC Status (RDP)            Active
```

```
  UncorrectedBlks(RDP)        315
  CorrectedBlks(RDP)          250
| Logins (FC)                 0
| SecCapableLogins (FC)       0
| EncryptedLogins (FC)        0
```

Report field definitions (with the `-metrics` parameter)

ID Specifies the fully qualified port ID.

Date
Specifies the current time stamp for the I/O port performance counters. For example, 08/11/05 02:23:49 is the way that this value is reported.

Byteread (FICON/ESCON)
Specifies the number of bytes that are read in increments of 128 KB.

Bytewrit (FICON/ESCON)
Specifies the number of bytes that are written in increments of 128 KB.

Reads (FICON/ESCON)
Specifies a value that is based on the extended count-key-data (ECKD) architecture data received operations.

Writes (FICON/ESCON)
Specifies a value that is based on the ECKD architecture data transferred operations.

Timeread (FICON/ESCON)
Specifies a value that is based on the ECKD data received (read-accumulated time) on a channel. The displayed value is based on increments of 16 milliseconds.

Timewrite (FICON/ESCON)
Specifies a value that is based on the ECKD data transferred (write-accumulated time) on a channel. The displayed value is based on increments of 16 milliseconds.

CmdRetries (FICON)
Specifies the number of retries that are requested for transport mode write operations because not enough buffers were available to receive unsolicited data.

Note: A dash (-) means that the data is not available.

TransferReady(FICON)
Specifies the number of transport mode operations from channels that require transfer ready and that were received from channels that support transport mode command retry operations.

Note: A dash (-) means that the data is not available.

Bytewrit (PPRC)
Specifies a value for the remote mirror and copy data transferred operation in increments of 128 KB.

Byteread (PPRC)
Specifies a value for the remote mirror and copy data received operations in increments of 128 KB.

Writes (PPRC)
Specifies a value for the remote mirror and copy data transferred operations.

Reads (PPRC)
Specifies a value for the remote mirror and copy data received operations.

Timewrite (PPRC)
Specifies a value that is based on the remote mirror and copy data transferred (write-accumulated) time on a channel. The displayed value is based on increments of 16 milliseconds.

Timeread (PPRC)
Specifies a value for the remote mirror and copy data received (read-accumulated) time on a channel. The displayed value is based on increments of 16 milliseconds.

Byteread (SCSI)
Specifies a value that is based on the SCSI data received operations. The displayed value is based on increments of 128 KB.

Bytewrit (SCSI)
Specifies a value that is based on the SCSI data transferred operations. The displayed value is based on increments of 128 KB.

Reads (SCSI)
Specifies a value that is based on the SCSI data received operations.

Writes (SCSI)
Specifies a value that is based on the SCSI data transferred operations.

Timeread (SCSI)
Specifies a value that is based on the SCSI data received (read-accumulated) time on a channel. The displayed value is based on increments of 16 milliseconds.

Timewrite (SCSI)
Specifies a value that is based on the SCSI data transferred (write-accumulated) time on a channel. The displayed value is based on increments of 16 milliseconds.

LinkFailErr (FC)
Specifies the total number of miscellaneous Fibre Channel link errors.

LossSyncErr (FC)
Specifies the number of loss of synchronization errors. These errors occur when there is a confirmed and a persistent synchronization loss on the Fibre Channel link.

LossSigErr (FC)
Specifies the number of times that a loss of signal was detected on the Fibre Channel link when a signal was previously detected.

PrimSeqErr (FC)
Specifies the number of primitive sequence protocol error counts where an unexpected primitive sequence was received.

InvTxWordErr (FC)
Specifies the number of times a bit error was detected. Examples of bit errors are a code violation, an invalid special code alignment, or a disparity error.

CRCErr (FC)
Specifies the number of times the CRC of a received frame is in error.

LRSent (FC)
Specifies the number of times the port that is changed from an active (AC) state to a Link Recovery (LR1) state.

LRRec (FC)
This count is the number of times the port that is changed from an active (AC) state to a Link Recovery (LR2) state.

IllegalFrame (FC)
Specifies the number of frames that violated the Fibre Channel protocol. One example of a violation is an invalid frame header. The violation occurs when the first frame of a data sequence is missing and a subsequent data frame is detected as illegal.

OutOrdData (FC)
Specifies the number of times that an out-of-order frame is detected. The frame is either missing from a data sequence or it is received beyond the sequence reassembly threshold of the port.

OutOrdACK (FC)
Specifies the number of times that an out-of-order ACK (ACKnowledgment field of the TCP protocol) frame is detected. The frame is either missing from a data sequence or it is received beyond the sequence reassembly threshold of the port.

DupFrame (FC)
Specifies the number of times a frame was received that was detected as previously processed.

InvRelOffset (FC)
Specifies the number of times that a frame was received with bad relative offset in the frame header.

SeqTimeout (FC)
Specifies the number of times the port detected a timeout condition after it received a sequence initiative for a Fibre Channel exchange.

BitErrRate (FC)
Specifies the number of bit error (transmission words that are not valid) bursts for the previous 5-minute counting window. This number is reset every 5 minutes.

RcvBufZero (FC)
Specifies the number of 1-second intervals that the receive buffer credit was zero.

SndBufZero (FC)
Specifies the number of 1-second intervals that the send buffer credit was zero.

RetQFullBusy (FC)
Specifies the number of times that the FCP port returned "queue full" or "busy" status.

ExchOverrun (FC)
Specifies the number of Fibre Channel exchanges that were lost due to overdriving the host adapter port.

ExchCntHigh (FC)
Specifies the number of times that the Fibre Channel exchange count crossed the High Threshold.

ExchRemAbort (FC)
Specifies the number of times that a port received an abort.

CurrentSpeed (FC)
Specifies the current operating link speed.

%UtilizeCPU (FC)
Specifies the percentage of the CPU that is being used. The percentage is followed by a space and either "Average" or "Dedicated". "Average" means that the Host Adapter CPUs are being used as a shared resource for the Host Adapter ports. "Dedicated" means that the Host Adapter CPUs are a dedicated resource for each Host Adapter port.

TxPower (RDP)
Specifies the measured coupled TX output power in dBm and uW. The maximum value is 18.1 dBm (6500 uW).

If the value 0 displays when the link state is active, the function is not supported by the current level of microcode.

Note: A dash (-) means that the data is not available.

RxPower (RDP)
Specifies the measured received optical power in dBm and uW. The maximum value is 18.1 dBm (6500 uW).

If the value 0 displays when the link state is active, the function is not supported by the current level of microcode.

Note: A dash (-) means that the data is not available.

TransceiverTemp (RDP)
Specifies the internally measured transceiver temperature. This value is reported in range -128 C to + 128 C.

If the value 0 displays when the link state is active, the function is not supported by the current level of microcode.

Note: A dash (-) means that the data is not available.

SupplyVolt(RDP)
Specifies the internally measured supply voltage. This value is reported in mV and the value range is 0 - 6550 mV.

If the value 0 displays when the link state is active, the function is not supported by the current level of microcode.

Note: A dash (-) means that the data is not available.

TxBias (RDP)
Specifies the measured transmitter laser bias current. The value is reported in mA and the value range is 0 - 131 mA.

If the value 0 displays when the link state is active, the function is not supported by the current level of microcode.

Note: A dash (-) means that the data is not available.

ConnectorType (RDP)
Specifies that the connector type displays one of the following values:

SFP+
Specifies a transmit power level of small form factor pluggable device.

Unknown
Specifies that the connector type might not be determined.

Dash (-)
Specifies that the data was not available.

TxType (RDP)
Specifies the port TX Type displays one of the following values:

Laser-SW
Specifies a short-wave laser port.

Laser LC 1310-LW
Specifies a long-wave laser LC 1310 nm port.

Laser LL 1550-LW
Specifies a long-wave laser LL 1550 nm port.

Non-Optical
Specifies that the port is not optical or is of another type.

Unknown
Specifies that the port type is not recognized.

Dash (-)
Specifies that the data was not available.

FECStatus(RDP)
Specifies the forward error correction status. One of the following values displays:

Active
Specifies that a forward error correction is active on the link.

Inactive
Specifies that a forward error correction is not active on the link.

UncorrectedBlks(RDP)
Specifies the number of bad data blocks that were not corrected by the forward error correction (FEC).

Note: A dash (-) means that the data is not available.

CorrectedBlks(RDP)
Specifies the number of bad data blocks that were corrected by forward error correction (FEC).

Note: A dash (-) means that the data is not available.

Security Login Counts
Specifies information about the security login counts. One of the following values is displayed:

ID The fully qualified port ID.

Logins
The number of remote hosts or peer *N*_ports current logged in to the specified I/O port.

SecCapableLogins
The number of logins that are capable of Fibre Channel security protocols. If the port is security-enabled, the remote ports are expected to perform authentication.

AuthLogins
The number of security capable logins that have negotiated link authentication.

EncryptedLogins
The number of remote host or peer *N*_ports that are currently logged into the specified port. These ports successfully authenticated and negotiated an encryption algorithm for the IBM Fibre Channel Endpoint Security function.

Example 3

Invoking the showioport command, with the -rdp parameter

```
dscli> showioport -dev IBM.1400-1B1 -rdp  I0103
ID                                I0103
WWPN                              50:05:07:63:03:03:00:45
Attached WWPN                     50:05:07:63:03:FF:C0:45
Physical Type                     FC-FS-3
Link Failure Error                0
Loss of sync Error                0
Loss of Signal Error              0
Primitive Sequence Error          0
Invalid Transmission Word Error   0
CRC Error                         0
FEC Status                        Active
Uncorrected Blocks                0
Corrected Blocks                  0
Port Speed Capabilities           4GFC 8GFC 16GFC (Configured to 16GFC)
Port Operating Speed              16GFC
Advertised B-B Credit             83
Attached Port B-B Credit          16
Nominal RTT Link Latency          5000 ns
ConnectorType                     SFP+
TxType                            Short Wave Laser
Transceiver Temperature           58.0 C        [Operating Range -128 - +128 C]
TxBiasCurrent                     14.0 mAmps    [Operating Range 0 - 131 mAmps]
Transceiver Supply Voltage        3327.0 mV     [Operating Range 0 - 3600mVolts]
RxPower                                         [Operating Range 0 - 6550 uW]
TxPower                                         [Operating Range 0 - 6550 uW]
Last SFP Read time                04/17/2017 20:17:44 CST
```

```
================================SFP Parameters Alarm Levels============================
Element                     High Warning     Low Warning    High Alarm    Low Alarm
=======================================================================================
Transceiver Temperature        0                0              0              0
Tx Bias Current                0                0              0              0
Transceiver Supply Voltage     0                0              0              0
Tx Power                       0                0              0              0
Rx Power                       0                0              0              0
==================================Attached Port========================================
ID                                N/A
WWPN                              50:05:07:63:03:FF:C0:45
Attached WWPN                     50:05:07:63:03:03:00:45
Physical Type                     FC-FS-3
Link Failure Error                0
Loss of Sync Error                0
Loss of Signal Error              0
Primitive Sequence Error          0
Invalid Transmission Word Error   0
CRC Error                         0
FEC Status                        Active
Uncorrected Blocks                0
Corrected Blocks                  0
Port Speed Capabilities           4GFC 8GFC 16GFC (Configured to 16GFC)
Port Operating Speed              16GFC
Advertised B-B Credit             83
Attached Port B-B Credit          16
Nominal RTT Link Latency          5000 ns
ConnectorType                     SFP+
TxType                            Short Wave Laser
Transceiver Temperature           58.0 C            [Operating Range -128 - +128 C]
TxBiasCurrent                     14.0 mAmps        [Operating Range 0 - 131 mAmps]
Transceiver Supply Voltage        3327.0 mV         [Operating Range 0 - 3600 mV]
RxPower                                             [Operating Range 0 - 6550 uW]
TxPower                                             [Operating Range 0 - 6550 uW]
Last SFP Read time                04/17/2017 20:17:44 CST

==============================SFP Parameters Alarm Levels==================================
Element                       High Warning     Low Warning    High Alarm    Low Alarm
===========================================================================================
Transceiver Temperature        0                0              0              0
Tx Bias Current                0                0              0              0
Transceiver Supply Voltage     0                0              0              0
Tx Power                       0                0              0              0
Rx Power                       0                0              0              0
```

Port Diagnostic Parameters (with the -rdp parameter)

ID Specifies the port ID. Or the fully qualified port ID in the form `manufacturer.type-model.serial number/portID`, when it is used with the **-fullid** parameter. For example, `I0201IBM.2107-75FA120/I0110`.

Note: The port ID component is four hexadecimal characters in the format `0xEEAP`, where `EE` is a port enclosure number (00-17 for 2107, 00-01 for 1750), `A` is the adapter number (1,2,4,5 for 2107, 0-3 for 1750), and `P` is the port number (0-3). The number is prefixed with the letter `I`.

WWPN
Specifies the worldwide port name of the port. For example, `50:05:07:63:03:03:00:45` is the way that this value is reported.

Note: A dash (-) means that the data is not available.

Attached WWPN
Specifies the worldwide port name of the attached port. For example, `50:05:07:63:03:FF:C0:45` is the way that this value is reported.

Note: A dash (-) means that the link is inactive, or the data is not available.

Physical Type
Specifies the physical type of the port, which includes one of the following types:

No Information
Specifies that no information about the physical type is provided.

FC-FS-3
Specifies that the sending VN_Port uses an FC-FS-3 PN_Port or PF_Port.

Lossless Ethernet MAC
Specifies that the sending VN_Port uses a lossless Ethernet MAC.

Unknown
Specifies unexpected values of physical type.

Note: A dash (-) means that the data is not available.

Link Failure Error
Specifies a link failure error count as a number value. This count is the number of miscellaneous Fibre Channel link errors. These errors might indicate that an unexpected NOS is received, or a link state machine failure is detected.

Note: A dash (-) means that the data is not available.

Loss of Sync Error
Specifies a loss of synchronization error count as a number value. This count is the number of loss of synchronization errors where it is a confirmed, persistent synchronization loss on the Fibre Channel link.

Note: A dash (-) means that the data is not available.

Loss of Signal Error
Specifies a loss of signal error count as a number value. This count is the number of times that a loss of signal was detected on the Fibre Channel link where a signal was previously detected.

Note: A dash (-) means that the data is not available.

Primitive Sequence Error
Specifies a primitive sequence error count as a number value. This count is the number of primitive sequence protocol error counts where an unexpected primitive sequence was received.

Note: A dash (-) means that the data is not available.

Invalid Transmission Word Error
Specifies an invalid transmission word error count as a number value. This count is the number of times that a *bit* error was detected. Examples of bit errors include a code violation, invalid special code alignment, or disparity errors.

Note: A dash (-) means that the data is not available.

CRC Error
Specifies a CRC error count as a number value. This count is the number of times a received frame CRC is in error.

Note: A dash (-) means that the data was not available.

FEC Status
Specifies the forward error correction (FEC) status. One of the following values is displayed:

Active
Specifies that a forward error correction is active on the link.

Inactive
Specifies that a forward error correction is not active on the link.

Note: A dash (-) means that the data is not available.

Uncorrected Blocks
Specifies the number of bad data blocks that were not able to be corrected.

Note: A dash (-) means that the data is not available.

Corrected Blocks
Specifies the number of bad data blocks that were corrected.

Note: A dash (-) means that the data is not available.

Port Speed Capabilities
Specifies the operating speed capabilities of the port. All possible port speed capabilities are listed. For example, `4GFC 8GFC 16GFC`. If this setting is configured, it indicates that the speed is configured to a fixed speed. For example, `4GFC 8GFC 16GFC (Configured to 16GFC)`.

Note: A dash (-) means that the data is not available.

Port Operating Speed
Specifies the current port operating speed, if it is set. If the link it not active, `Inactive` is displayed.

Note: A dash (-) means that the data is not available.

Advertised B-B Credit
Specifies the FC_Port Buffer-to-Buffer Credit as a number value. This value is the advertised buffer-to-buffer credit for the Fx_Port or Nx_Port for which diagnostic parameters are being processed.

Note: A dash (-) means that the data is not available.

Attached Port B-B Credit
Specifies the attached FC_Port Buffer-to-Buffer Credit as a number value. This value is the advertised buffer-to-buffer credit that is received from the directly attached Fx_Port or Nx_Port.

Note: A dash (-) means that the link is inactive or the data is not available.

Nominal RTT Link Latency
Specifies the Nominal FC_Port Round Trip Time (RTT). This value is the nominal RTT of the link between the FC_Port and the attached FC_Port as determined by the FC_Port, where a value of zero indicates that the RTT is unknown. A nonzero value specifies the RTT of the link in nanoseconds.

Note: A dash (-) means that the link is inactive or the data is not available.

Connector Type
Specifies the type of connector. The value is either:

SFP+
Specifies that the connector type transmits the power level of a small form-factor pluggable (SFP) device.

Unknown
Specifies that the connector type cannot be determined.

Note: A dash (-) means that the data is not available.

Tx Type
Specifies the port Tx type, which is one of the following values:

Not Optical or Other Type
Specifies that the port is not optical or is of another type.

Short Wave Laser
Specifies a short-wave laser port.

Long Wave Laser LC 1310nm
Specifies a long-wave laser LC 1310 nm port.

Long Wave Laser LL 1530nm
Specifies a long-wave laser LL 1530 nm port.

Unknown
Specifies that the port type is not recognized.

Dash (-)
Specifies that the data was not available.

Transceiver Temperature
Specifies the internally measured temperature of the transceiver. The temperature range is -128 C to + 128 C.

Note: A dash (-) means that the data is not available.

Tx Bias Current
Specifies the measured current of the transmitter (Tx) laser bias. The range is 0 - 131 mAmps.

Note: A dash (-) means that the data is not available.

Transceiver Supply Voltage
Specifies the internally measured supply voltage of the transceiver. The range is 0 - 6.55 volts.

Note: A dash (-) means that the data is not available.

Rx Power
Specifies the Rx power. The range is 0 - 6.5 mW.

Note: A dash (-) means that the data is not available.

Tx Power
Specifies the Tx power. The range is 0 - 6.5 mW.

Note: A dash (-) means that the data is not available.

Last SFP Read time
Specifies the last time that the SFP was read. For example, `07-10-2015 UTC Fri 23:16:18`.

Note: A dash (-) means that the data is not available.

Remember: The same port diagnostic parameters also apply to the attached port.

SFP Alarm Levels (with the `-rdp` parameter)

Element
Specifies the SFP parameter name, which is one of the following values:

- Transceiver Temperature
- Tx Bias Current
- Transceiver Supply Voltage

- Tx Power
- Rx Power

High Warning
Specifies the high warning status. The possible values are:

0 Specifies that the high warning flag is not set.

1 Specifies that the high warning flag is set.

Dash (-)
Specifies that the diagnostic parameters are not valid. The response does not include valid values for SFP diagnostic parameters.

Low Warning
Specifies the low warning status. The possible values are:

0 Specifies that the low warning flag is not set.

1 Specifies that the low warning flag is set.

Dash (-)
Specifies that the diagnostic parameters are not valid. The response does not include valid values for SFP diagnostic parameters.

High Alarm
Specifies the high alarm status. The possible values are:

0 Specifies that the high alarm flag is not set.

1 Specifies that the high alarm flag is set.

Dash (-)
Specifies that the diagnostic parameters are not valid. The response does not include valid values for SFP diagnostic parameters.

Low Alarm
Specifies the low alarm status. The possible values are:

0 Specifies that the low alarm flag is not set.

1 Specifies that the low alarm flag is set.

Dash (-)
Specifies that the diagnostic parameters are not valid. The response does not include valid values for SFP diagnostic parameters.

Remember: The same SFP alarm levels also apply to the attached port.

Host commands

Various commands are available to help configure host connections and to display host connection information.

The following host connection commands are available:

chhost Modifies the configuration of a host.

chhostconnect
Modifies a SCSI host port configuration. You must ensure that the host port is offline to the host system before you process the **chhostconnect** command.

lshost Displays a list of hosts for a specified storage system and status information for each host in the list.

lshostconnect
Displays a list of host connections for a storage image and the status information for each host connection in the list.

lshosttype
Displays a list of known hosts, their associated port profiles, address discovery, and logical block size values.

lshostvol
Displays the mapping of host device names or volume names to machine type 2107 /242x volume IDs.

lsportprof
Displays a list of port profiles that are supported on a storage system and their recommended address discovery and logical block size values. This command is helpful to obtain the recommended values for the **mkhostconnect** command.

managehostconnect
Modifies the volume group assignment for a SCSI host port. Ensure that the host port is offline to the host system before you process the **managehostconnect** command.

mkhost Configures an open systems host system.

mkhostconnect
Configures the open systems hosts port attachments to Fibre Channel ports that are configured for FC-AL or SCSI-FCP topology.

rmhost Removes a host from a storage image.

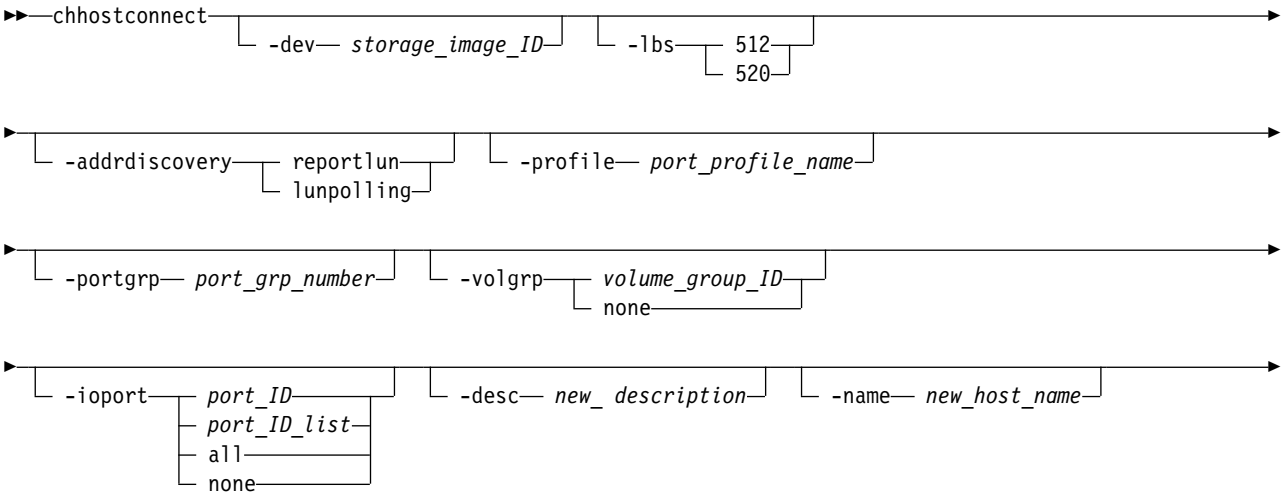
rmhostconnect
Removes a SCSI host port connection from a storage image.

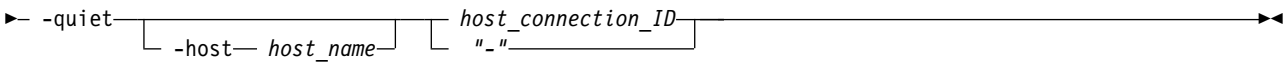
showhost
Displays detailed properties of an individual host.

showhostconnect
Displays the detailed properties of a specified storage image host connection.

chhostconnect

The **chhostconnect** command modifies a SCSI host port configuration.





Parameters

Notes:

1. The **chhostconnect** command can be disruptive to host system I/O operations if the affected host port is logged in to the target storage system. You must ensure that the host port is offline to the host system before you process the **chhostconnect** command.
2. Using the **-hosttype** parameter when you enter this command saves input and processing time. The **-hosttype** parameter supplies the same information as if you used the following three parameters:
   - **-profile**
   - **-addrdiscovery**
   - **-lbs**
3. If you are using the HP-UX operating system, see the volume restriction that is described under the **-addrdiscovery** parameter.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host connection ID. It is also required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

An example of a fully qualified storage image ID includes IBM.2107-75FA120.

-lbs 512 | 520
(Optional) Specifies the logical block size that is used by the host system. The block size must be compatible with the volume group type and the volume type configurations that apply to the host port connection. The 520 logical block size is typically used by the IBM System i models (OS/400).

Notes:

1. You cannot use the **-lbs** parameter and **-hosttype** parameter together, but you can use each one separately.
2. If you do not use the **-hosttype** parameter, use the **lsportprof** command to determine the block size that you need to specify for the **-lbs** parameter.

-addrdiscovery reportlun | lunpolling
(Optional) Specifies the method for identifying logical unit number (LUN) addresses.

- The **reportlun** method specifies that the host system can access up to 64 000 LUNs.
- The **lunpolling** method specifies that the host system can access up to 256 LUNs.

Notes:

1. You cannot use the **-addrdiscovery** parameter and **-hosttype** parameter together, but you can use each one separately.
2. For HP-UX operating systems, the number of volumes in the volume group must not exceed seven volumes. This restriction applies only when the **-addrdiscovery** parameter is set to *reportlun* and the associated volume group is of type **scsimap256**.

-profile *port_profile_name*
(Optional) Specifies the name of the host connection behavior profile. If the name includes a blank space, enclose the name with double quotation marks. For example, -profile “IBM pSeries - Sun”.

Notes:

1. You cannot use the **-profile** parameter and the **-hosttype** parameter together, but you can use each one separately.
2. If you do not use the **-hosttype** parameter, use the **lsportprof** command to obtain a list of available profiles.

-portgrp *port_grp_number*
(Optional) Specifies a user-assigned number that associates two or more host ports with access to a common volume group. Port group zero is reserved for ports that are not associated with a port group.

-volgrp *volume_group_ID* | **none**
(Optional) Specifies an available volume group or no volume group if the **none** subparameter is used. This command accepts a fully qualified volume group ID including the storage image ID or a shortened version if the **-dev** parameter is specified. The shortened version is a four-digit decimal number with no leading zeros, prefixed with the letter *V*.

A host connection can use only one volume group per storage image. (A single WWPN can access only one volume group per storage image.) Host operations cannot be initiated until a volume group ID is assigned.

If *none* is specified, the volume group ID assignment is removed from a SCSI host port object.

-ioport *port_ID port_ID_list* |**all**|**none**
(Optional) Specifies all, none, or, one or more I/O port IDs that allow host connection access to volumes. This command accepts a fully qualified port ID including the storage image ID or a shortened version if the **-dev** parameter is specified.

port_ID port_ID_list
Specifies that you can designate up to 128 ports for an open systems host attachment assignment. If you enter a list of I/O port IDs, access from the specified host connection to the specified volume group is allowed by using only the designated list of port IDs.

To specify a range of port IDs, separate the port IDs with a hyphen. You can list one or more port IDs or ranges of port IDs separated by commas with no spaces.

all Specifies that you want to add all I/O port IDs. This setting allows the specified host connection access to the designated volume group through all the associated I/O port IDs.

none Specifies that you do not want to add any I/O ports. If you do not specify I/O ports, the storage system is configured to allow host connection access to the specified volume group. The access is made by using any I/O port that is configured for SCSI-FCP topology.

Examples of -dev parameter use

If you specify the **-dev** parameter, you can use the shortened version of the **-ioport** parameter as follows:

```
dscli> chhostconnect -dev IBM.2107-75FA120 -ioport I0222 1
```

Where *1* represents the required parameter, **host_connection_ID**.

If you do not specify the **-dev** parameter and you specify the **-ioport** parameter, you must use the fully qualified version of the port ID with the **-ioport** parameter specified as follows:

```
dscli> chhostconnect -ioport IBM.2107-75FA120/I0222
 IBM.2107-75FA120/1
```

Where *IBM.2107-75FA120/1* represents the required parameter, **host_connection_ID**.

A port ID is prefixed with the letter *I* and consists of four hexadecimal characters in the format *EEAP*, where:

EE The I/O port enclosure number in the range of 00 - 17.

A The adapter number and is specified as 1, 2, 4, or 5.

P The port number (0 - 3).

Note: Changing the I/O port values can result in a disruption of current logins by the host systems.

-desc *new_description*
(Optional) Specifies the description that you defined for the SCSI host port. The description is limited to 256 byte or 128 double-byte characters.

-name *new_host_name*
(Optional) Specifies the name that you assigned for host system or port. The name is limited to 32 characters.

-quiet
(Optional) Turns off the confirmation prompt for this command.

-host *host_name*
(Optional) Specifies the name of the host to which the host connection is assigned.

host_connection_ID | **-**
(Required) Specifies the host connection ID, which is a unique identifier that uses any number from 0 - FFFE within the scope of a storage image. This parameter accepts a fully qualified ID or a shortened version if the **-dev** parameter is specified.

If you do not specify the **-dev** parameter and you specify the **host_connection_ID** parameter, you must use the fully qualified version of the host connection ID as follows:

```
dscli> chhostconnect -desc newdescription IBM.2107-75FA120/1
```

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Attention: Use caution when you work with connection IDs to ensure that you specified the correct connection that you want to change. For example, if you intend to change connection ID 0005 and type 000, the system changes connection ID 0. Or, if you want to change connection ID 0020 and type 002, the system changes connection ID 2. The system does not recognize the leading zeros, and 000 is interpreted as connection ID 0 and 002 is interpreted as connection ID 2.

Invoking the chhostconnect command

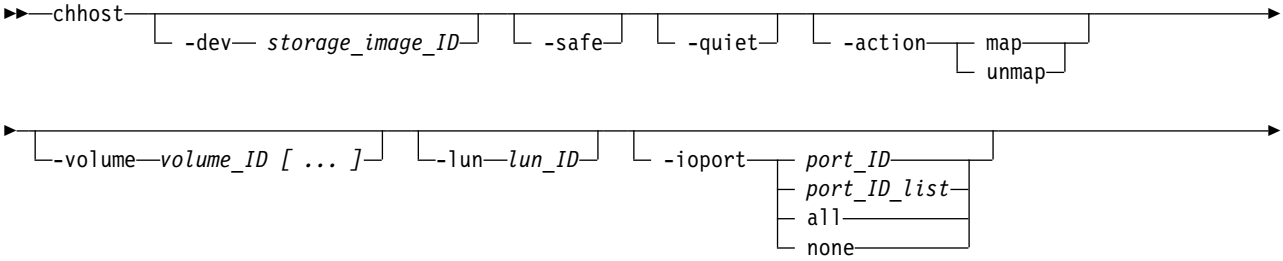
```
dscli> chhostconnect -dev IBM.2107-75FA120 -name host_1_port_2 1
```

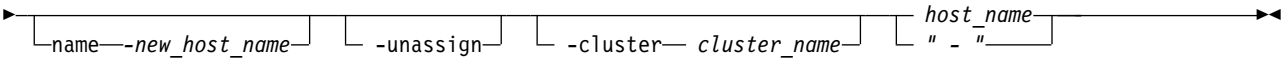
The resulting output

```
Host connection 1 successfully modified.
```

chhost

The **chhost** command modifies the configuration of a host.





Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified port ID. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-safe
(Optional) Enables checks for the safe removal of the volume from the host.

-quiet
(Optional) When volumes are mapped or unmapped from a host, and if more than one host share volume mappings, a prompt displays. You can confirm whether (Yes or No) to proceed mapping or unmapping the volumes from a host. This parameter turns off the confirmation prompt.

-action *map unmap*
(Optional) Specifies the type of action for volumes.

map
Specifies that the volumes are added to the host.

unmap
Specifies that the volumes are removed from the host.

Note: This **chhost** command fails if you specify the **-volume** without the **-action** parameter. You cannot specify the **-action** parameter with the **-cluster** or **-unassigned** parameters.

-volume *volume_ID [...]*
(Optional) Specifies an array of one or more volume IDs or volume ID ranges to be added or removed from the specified host. A volume ID range is defined by two volume IDs that are separated by a hyphen. Multiple volume IDs or volume ID ranges must be separated by comma. The volume IDs are entered as hexadecimal numbers in the range 0x0000-0xFEFF.

You must separate multiple volume IDs or ranges of volume IDs with a comma between each ID or range of IDs.

Notes:

1. For SCSI MAP 256, the array or ranges cannot exceed 256 volume ID entries. Otherwise, up to 64,384 entries are allowed.
2. The **chhost** command fails if you specify the **-volume** parameter and do not specify the **-action** parameter.

-lun *lun_ID*
(Optional) Specifies the LUN ID in hexadecimal value (00 - FF), which is mapped to the specified volume ID for an SCSI-MAP256 type host. If multiple volume IDs are specified by the **-volume** parameter, the LUN IDs are assigned consecutively in incremental order.

This parameter is only valid for a SCSI-MAP 256 type host. If you specify this parameter for any other type of host, the command fails.

-ioport *port_ID port_ID_list*|**all**|**none**
(Optional) Specifies all, none, one, or more I/O port IDs that allow host connection access to volumes.

all
Specifies that you want to add all I/O port IDs to allow the specified host connection access to the designated volume group through all the associated I/O port IDs.

none
Specifies that you do not want to add any I/O ports. If you do not specify I/O ports, then the storage system is configured to allow host connection access to the specified volume group. The access is through any I/O port that is configured for SCSI-FCP topology.

port_ID port_ID_list
Specifies that you can designate up to 128 ports for an open systems host attachment assignment. You must separate multiple port IDs or ranges of port IDs with a comma between each ID or range of IDs. To specify a range of port IDs, separate the port IDs with a hyphen.

If you enter a list of I/O port IDs, access from the specified host connection to the specified volume group is allowed by using only the designated list of port IDs

A port ID is four hexadecimal characters in the format *EEAP*, where:

EE The I/O port enclosure number in the range of 00 - 17.

A The adapter number and is specified as 1, 2, 4, or 5.

P The port number (0 - 3).

This number is prefixed with the letter *I*.

-name *-new_host_name*
(Required) Specifies the new host system name that you assign, limited to 32 byte, or 16 double byte characters. If the name contains a blank space, you must surround the name in double quotation marks.

-unassign
(Optional) Unassign the host from the cluster. You cannot use this parameter with **-cluster** or **-action** parameters.

-cluster *cluster_name*
(Optional) Specifies the cluster name.

Note: If the cluster name contains a blank space, you must surround the name in double quotation marks. You cannot use this option with -**unassign** or **-action** parameters.

host_name | **-**
(Required) Specifies the name of the host system to be modified.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the chhost command

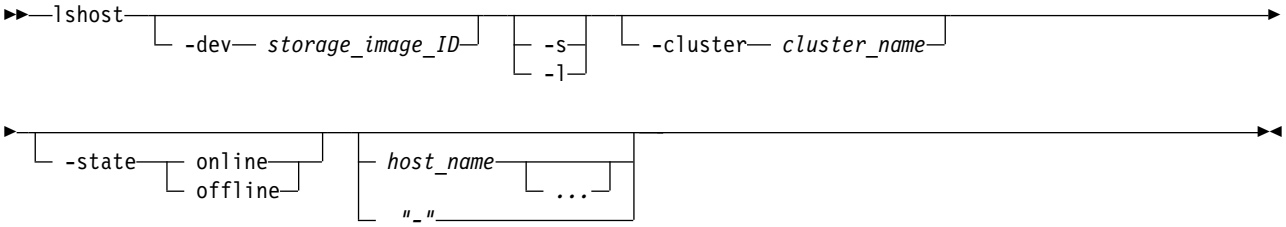
```
dscli> chhost -name host_1_port_2 host_1
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2016 DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
Host host_1 successfully modified.
```

lshost

The **lshost** command displays a list of hosts for a specified storage system and status information for each host in the list.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host connection ID. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-s
(Optional) Displays only the host name in the command output. You cannot specify the **-s** and **-l** parameters together.

-l
(Optional) Displays the default output plus the description that you included for each host in the list. You cannot specify the **-s** and **-l** parameters together.

-cluster *cluster_name*
(Optional) Displays all hosts under this cluster.

-state online | offline
(Optional) Displays all hosts with the specified state of online or offline.

host_name ... | -
(Optional) Displays host information for the host names that are specified. Multiple names must be separated with a white space between each name or ID. The ellipsis (...) indicates that you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the lshost command

```
dscli> lshost -l
```

The resulting output

Name	Type	State	Cluster	AddrMode	Volumes	Host ports	I/O ports
Host_1	VMWare	Offline	cluster_1	scsimap256	8	8	32
Host_2	VMWare	Offline	-	scsimap2566	6	8	32
Host_3	VMWare	Offline	-	scsimap2560	0	8	32

Report field definitions

Name*
Specifies the name that you assigned to the host.

Type
Specifies the name of the host type as defined by the operating system.

Unknown displays when the information is not available and indicates that the host type was not specified when the host connection was created or modified.

State
Specifies one of the states of the specified host.

online
If at least one host port is in the logged-in state, the state displays as online.

offline
If no host ports were added to the specified host, or if all of the added host ports are logged out, the state displays as offline.

Cluster
Specifies the name of the cluster that the specified host is assigned to.

AddrMode+
Specifies the method that is used by the host to discover the LUNs that are available to its host ports (for example, scsimask, scsimap256, os400mask).

Volumes+
Specifies the number of volumes that are mapped to the specified host.

Host ports+
Specifies the number of host ports that are mapped to the specified host.

I/O ports+
Specifies the number of I/O ports that the specified host is allowed to log in to. A value of *All* indicates that the specified SCSI host is allowed to log in to any I/O ports.

Key:

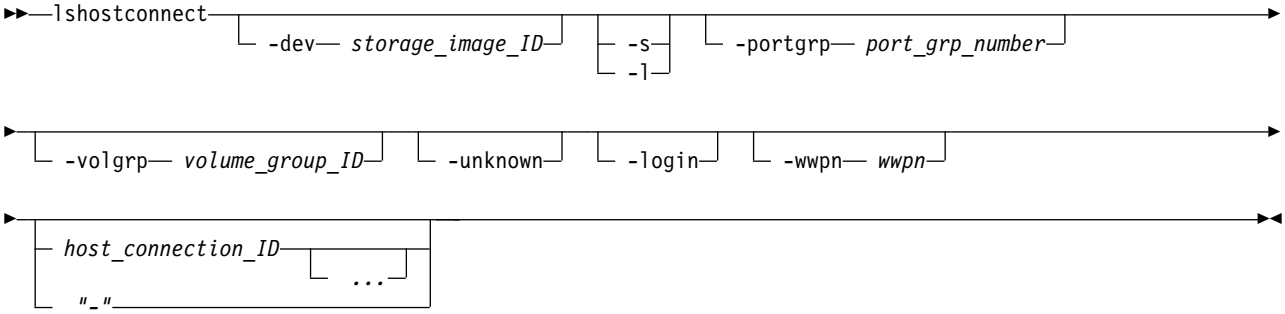
* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

lshostconnect

The **lshostconnect** command displays a list of host connections for a storage image and the status information for each host connection in the list.

You can also use this command to obtain a list of worldwide port numbers (WWPNs) from a system-detected-unknown host port. You can use these WWPNs to create a host connection by using the **mkhostconnect** command.



Parameters

-dev *storage_image_ID*
(Optional) Displays the host connections for the specified storage image. A storage image ID consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not

specify a fully qualified host connection ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

An example of a fully qualified storage image ID: IBM.2107-75FA120.

-s
(Optional) Specifies the host connection IDs. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Specifies the default output and your description for each host connection in the list. You cannot use the **-l** and the **-s** parameters together.

-portgrp *port_grp_number*
(Optional) Specifies the host connections that have an associated group number.

Note: You cannot use the **-portgrp** parameter with the **-unknown** or **-login** parameters.

-volgrp *volume_group_ID*
(Optional) Specifies that only the host connections with the specified volume group ID are to be displayed. The volume group ID is a four-digit decimal number with no leading zeros, prefixed with the letter *V*.

Note: You cannot use the **-volgrp** parameter with the **-unknown** or **-login** parameters.

-unknown
(Optional) Specifies that a list of logged in host ports (WWPNs), that are not recognized as being associated with the designated storage system, be displayed.

Note: The list of logged in host ports includes all of the host ports that the storage system detects. The list does not account for changes that the storage system did not detect. For example, the storage system cannot detect that a cable is disconnected from the port of the host device or that a fabric zoning change occurred. In these cases, the host might not be able to communicate with the storage device anymore. However, the storage device might not detect this condition and still views the host as logged in.
This parameter generates a list report that contains the following three information fields:

- WWNN
- WWPN
- ESSIOport

Note: You cannot use the **-unknown** parameter with the **-portgrp**, **-volgrp**, **-login**, or *host_connection_ID* parameters.

-login
(Optional) Specifies that a list is to be displayed of host port (WWPNs) that are logged in and sorted by the ESS I/O port IDs for known connections. The report displays one line of information per connection. However, no information is displayed for a FICON connection.

Notes:

1. Known logins are logins that you created for host connection, and Remote Mirror and Copy paths and anonymous connections.
2. You cannot use the **-login** parameter with the **-unknown**, **-portgrp**, **-volgrp**, or *host_connection_ID* parameters.

-wwpn *wwpn*
(Optional) Specifies that only the host connect objects for the specified WWPN are displayed.

host_connection_ID ... | **-**
(Optional) Specifies that host connection information for the specified host connection IDs be

displayed. This parameter accepts a fully qualified ID (includes *manufacture. machine type, serial number/hostconnectID*) or a shortened version if the **-dev** parameter is specified.

Note: You cannot use the *host_connection_ID* parameter with the **-login** or **-unknown** parameters.

The ellipsis (...) indicates that, optionally, you can specify multiple values.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

An example of a fully qualified host connection ID: IBM.2107-75FA120/2.

Example of a shortened version host connection ID: 0002

Example

Note: You can receive different reports when you use the **lshostconnect** command, one for the **-unknown** parameter, one for the **-login** parameter, one for the **-l** parameter, and one for the **-s** parameter. The reports that are associated with the **-unknown**, **-login**, and **-l** parameters are provided in this description.

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output reports that are associated with the **lshostconnect** command.

Invoking the lshostconnect command without the -unknown parameter

```
dscli> lshostconnect -dev IBM.2107 75FA120 -l
```

The resulting output The following table displays four ports, which is possible when you query a DS8000 model.

Profile	portgrp	volgrpID	atchtopo	ESSIO port	Speed	Desc	host
IBM pSeries - AIX	0	V100	-	I0111,I0121 I0211,I0221	Unknown	SCSI1	Production01
IBM pSeries - AIX	0	V100	-	All	Unknown	SCSI2	Production
IBM pSeries - pLinux	0	V100	-	I0111,I0121 I0211,I0221	Unknown	SCSI3	Production01
IBM pSeries - pLinux	0	V100	-	-	Unknown	SCSI4	Production01

Example

Invoking the lshostconnect command with the -unknown parameter

```
dscli> lshostconnect -dev IBM.2107 75FA120 -unknown
```

The resulting output

WWNN	WWPN	ESSIOport
3007ACF30A239900	3007ACF30A2399E0	I0111, I0121, I0211, I0221
3007ACF30A239900	3007ACF30A2399E1	I0121
3007ACF30A239900	3007ACF30A2399E2	I0111, I0121, I0211, I0221
3007ACF30A239900	3007ACF30A2399E3	I0111

Example

Invoking the lshostconnect command with the -login parameter

```
dscli> lshostconnect -dev IBM.2107-75FA120 -login
```

The resulting output

WWNN	WWPN	ESSIOport	LoginType	Name	ID
3007ACF3 0A239900	3007ACF3 0A2399E0	I0111	SCSI	MyHostA	1
3007ACF3 0A239900	3007ACF3 0A2399E1	I0111	SCSI	MyHostB	1
3007ACF3 0A239900	3007ACF3 0A2399E2	I0221	SCSI	-	-

Report field definitions when the -unknown or -login parameter is not used

Name
Host connection/SCSI port nickname.

The name is limited to 32-byte or 16 double-byte characters.

ID A fully qualified host connection ID: *manufacturer.type-serial number/hostconnectID*.

The host connection ID component is a unique identifier (0 - FFFE) within the scope of a storage system.

WWPN
Indicates the worldwide port name (WWPN) for this host system port.

HostType
Indicates the name of the host type.

Unknown is displayed when the information is not available and indicates that the host type was not specified when the host connection was created or modified.

LBS
Indicates the logical block size that is used by the designated host system and host system port.

The logical block setting must be compatible with the volume group type that is configured for volume access. The 520 block size is typically used for IBM i host system attachments.

Addrdiscovery
Indicates the LUN address discovery method that is used by the designated host system and host system port.

The LUN Discovery method must be compatible with the volume group type that is configured for volume access.

The Poll LUNs method enables access to a maximum of 256 LUNs. The Report LUNs method enables access to a maximum of 64 000 LUNs.

Profile
Indicates the name of the host connection behavior profile.

Portgrp
Indicates the host port group ID. This ID ties together a group of SCSI host port objects that are accessing a common volume group. If the port group value is set to zero, the host port is not associated with any port group.

VolgrpID
Indicates the volume group ID. This ID is a unique identifier within the DS8000 for the SCSI volume group that the specified SCSI host port is assigned to.

Atchtopo
The atchtopo is an attribute of the I/O port and, under certain conditions, might display a value that is inconsistent with the value reported by the **lsioport** command. Because this inconsistency cannot be corrected in all cases, the atchtopo attribute always displays as a " - " value. To obtain the I/O port topology, use the **lsioport** or **showioport** commands.

ESSIOport
Indicates the array of port IDs that the designated host port is logged in to.

A port ID is prefixed with the letter "I" and consists of four hexadecimal characters in the format *EEAP*, where:

- *EE* is an I/O port enclosure number in the range of 00 - 17.
- *A* is the adapter number and is specified as 1, 2, 4, or 5.
- *P* is the port number (0 - 3).

Speed
Speed is an attribute of the I/O port and, under certain conditions, might display a value that is inconsistent with the value reported by the **lsioport** command. Because this inconsistency cannot be corrected in all cases, the speed attribute always displays a value of "Unknown". To query the speed of an I/O port, use the **lsioport** or **showioport** commands.

Desc
Indicates the description that you defined for the SCSI host port. The description is limited to 256 byte or 128 double-byte characters.

host
Indicates the host name that is used in the Storage Management GUI, which is different from the nickname that is displayed in the **Name** column. If a name is not available, a value of "-" displays.

Report field definitions when the -unknown parameter is used

WWNN
Indicates the worldwide node name for the designated host system.

WWPN
Indicates the worldwide port name for the designated host system port.

ESSIOport
Indicates the array of port IDs that the designated host port is logged in to.

A port ID is prefixed with the letter "I" and consists of four hexadecimal characters in the format *EEAP*, where:

- *EE* is an I/O port enclosure number in the range of 00 - 17.
- *A* is the adapter number and is specified as 1, 2, 4, or 5.
- *P* is the port number (0 - 3).

Report field definitions when the -login parameter is used

WWNN
Indicates the worldwide node name (WWNN) for this host system.

WWPN
Indicates the worldwide port name (WWPN) for this host system port.

ESSIOport
Indicates the port ID that the designated host port is logged in to.

A port ID is prefixed with the letter "I" and consists of four hexadecimal characters in the format *EEAP*, where:

- *EE* is an I/O port enclosure number in the range of 00 - 17.
- *A* is the adapter number and is specified as 1, 2, 4, or 5.
- *P* is the port number (0 - 3).

LoginType
Indicates the type of login such as SCSI.

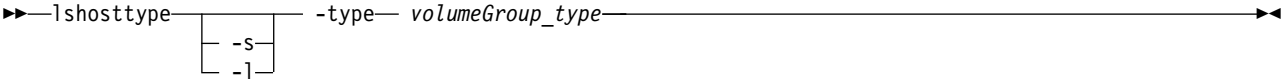
Name
Indicates the name that you assigned to the host. If a name is not assigned, a " - " value is displayed.

ID A fully qualified host connection ID: *manufacturer, machine type, serial number/hostconnectID*.

The host connection ID component is a unique identifier (0 - FFFE) within the scope of a storage system.

lshosttype

The **lshosttype** command displays a list of known hosts, their associated port profiles, address discovery, and logical block size values. Use this command to get the available host types for the **mkhostconnect** command.



Parameters

-s
(Optional) Displays the host types only. You cannot use the **-l** and **-s** parameters together.

-l
(Optional) Displays the default output for the specified host type. You cannot use the **-l** and **-s** parameters together.

-type *volumeGroup_type*
(Required) Displays only those host types that are associated with the specified volume group type.

volumeGroup_type
Only one type can be queried at a time. The following list provides the choices that can be specified.

- ficonall
- scsiall
- scsimask
- scsimap256
- os400all
- os400mask

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lshosttype** command.

Invoking the lshosttype command

```
dscli> lshosttype -l -type scsiall
```

The resulting output

Host Type	Profile	AddrDiscovery	LBS	Description
pSeries	IBM pSeries - AIX	reportlun	512	IBM pSeries, RS/6000 and RS/6000 SP Servers (AIX)
zLinux	IBM zSeries - zLinux	lunpolling	512	IBM zSeries Servers (Linux)

Report column definitions

Host Type*
Specifies the name of the specific host type.

Profile
Specifies the name of the host connection behavior profile. The port profile specifies a host or operating system type.

AddrDiscovery
Specifies the address discovery method. One of the following values is displayed:

LUN Polling
Specifies that host system LUN access is limited to a maximum of 256 LUNs.

Report LUN
Specifies that host system LUN access is limited to a maximum of 64 K LUNs.

LBS
Specifies the logical block size. One of the following values is displayed:

- 512 - This value is displayed for all hosts except OS400.
- 520 - This value is displayed for an OS400 host.

Description
Specifies more host type details.

Key:

* Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

lshostvol

The **lshostvol** command displays the mapping of host device names or volume names to machine type 2107 volume IDs. (This command is not supported on the i5/OS).

```
►►— lshostvol ————————————————————————————————————————————————————————————►◄
```

Parameters

This command has no parameters.

Notes:

1. The **lshostvol** command displays only volumes that are accessed by using a direct Fibre Channel path when you use the command on an OpenVMS host system that is a member of an OpenVMS cluster. The command output does *not* display information about the following OpenVMS cluster devices:

- Volumes to which the host system has only Mass Storage Control Protocol (MSCP) paths.

  Note: Mass Storage Control Protocol uses two queues. Into one queue, packets are placed which fully describe the commands to be run by the mass storage subsystem. To start an I/O request, the creates a small data structure in memory and appends it to a "send" queue. If this packet is the first packet in the send queue, it wakes the MSCP controller. After the command is processed, an appropriate status packet is placed into the second queue to be read by the CPU.
- Volumes to which the host system uses only MSCP paths now even though it has both MSCP and direct paths.

2. If you did not install the IBM Multipath Subsystem Device Driver (SDD), the virtual path (vPath) name is not displayed.
3. On a Red Hat Enterprise Linux system, attached devices might be detected by the HBA driver, but they are not registered with the operating system. Normally, the operating system is set up to automatically detect all LUNS. However, if this detection does not occur automatically, you must issue the following command for every volume (LUN):

   ```
   echo
   scsi add-single-device host# channel# lun# >/proc/scsi/scsi
   ```

   If SDD is installed on your system, you can run the scsiscan script to detect all the LUNs.
4. If the user that is running the DS CLI on the host does not have permissions to view the host volumes, the **lshostvol** command returns no host volumes found.

Example

The information that is displayed on the report that is generated from the **lshostvol** command is different depending on whether you installed SDD. The following example tables indicate the differences.

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output reports that are associated with the **lshostvol** command.

Invoking the lshostvol command

```
dscli> lshostvol
```

The resulting output with SDD installed

Disk Name	Volume ID	Vpath Name
my_vol_01,my_vol_04	IBM.2107-75DA110/175D	vpath01
my_vol_02,my_vol_05	IBM.2107-75EA120/175E	vpath02
my_vol_03,my_vol_06	IBM.-75FA130/175F	vpath03

The resulting output without SDD installed

Device/Volume Name	Volume ID	Vpath Name
Hdisk01	IBM.2107-75DA110/175D	-
Hdisk02	IBM.2107-75EA120/175E	-
Hdisk03	IBM.2107-75FA130/175F	-

Report column definitions

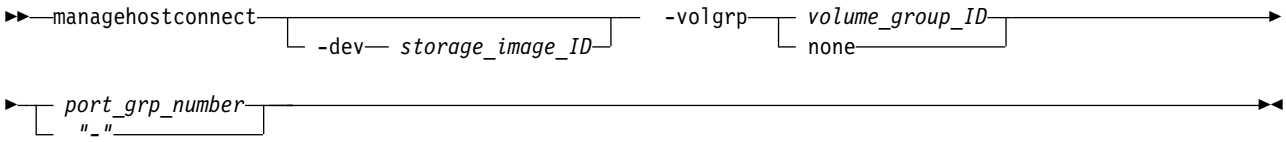
Device/Volume name
Specifies the nickname that you assigned to the device or volume. When SDD is installed, this column reports the volume name instead of the device name.

Volume ID
Specifies the ID of the storage system.

Vpath name
Specifies the virtual path name. When SDD is not installed, this value is reported as " - ".

managehostconnect

The **managehostconnect** command modifies the volume group assignment for a SCSI host port.



Parameters

Notes:

1. The **managehostconnect** command can be disruptive to host system I/O operations if the affected host port is logged on to the target storage unit. Ensure that the host port is offline to the host system before you process the **managehostconnect** command.
2. This command is used more effectively after you issue the **lshostconnect** or **showhostconnect** commands and analyze the reports that are generated by these commands. The information that is reported by these commands can help you ensure that you specify the correct port group number when you issue the **managehostconnect** command.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-volgrp *volume_group_ID* | **none**
(Required) Specifies that the SCSI host port connections that are associated with the specified port group number are modified to access this volume group ID. A volume group ID is a four-digit decimal number with no leading zeros, prefixed with the letter **V**.

If *none* is specified, the volume group ID assignment is removed from all SCSI host port objects that are associated with a common port group number.

Example: -volgrp none.

port_grp_number | **-**
(Required) Specifies the SCSI host port group number that associates two or more host ports as having access to a common volume group.

If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the managehostconnect command

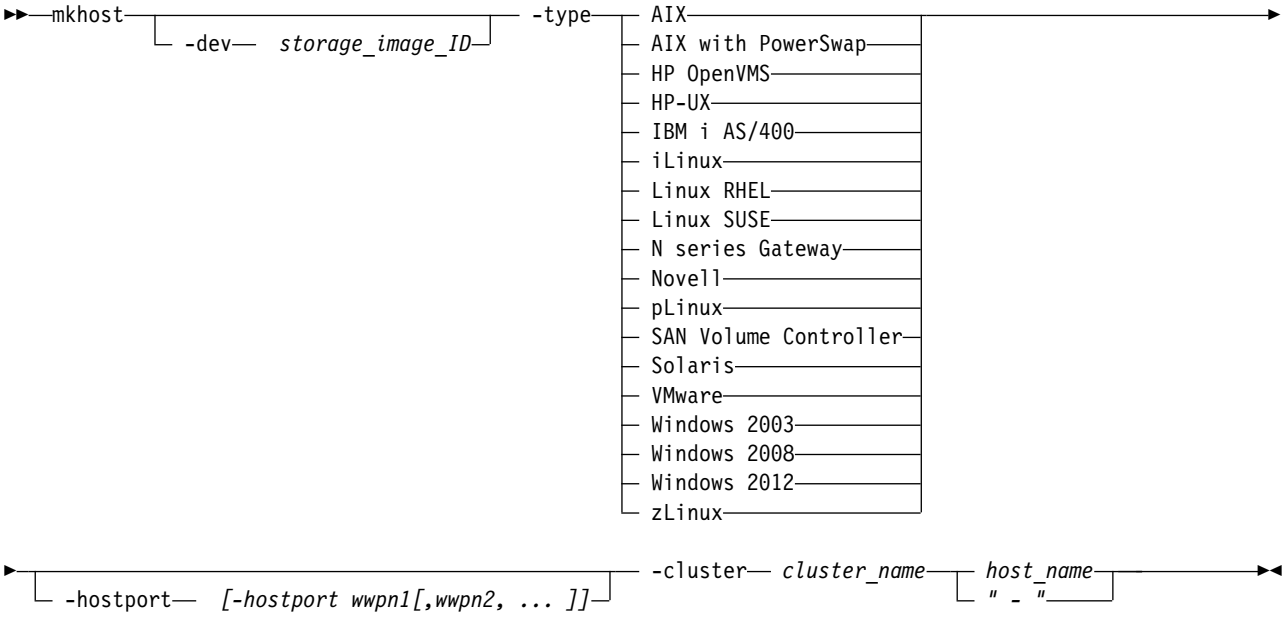
```
dscli> managehostconnect -dev IBM.2107-75FA120 -volgrp 11 1
```

The resulting output

```
Port group number 1 successfully modified.
```

mkhost

The **mkhost** command configures an open systems host system.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host connection ID. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-type AIX | AIX with PowerSwap | HP OpenVMS | HP-UX | IBMi AS/400 | iLinux | Linux RHEL | Linux SUSE | N series Gateway | Novell | pLinux | SAN Volume Controller | Solaris | VMware | Windows 2003 | Windows 2008 | Windows 2012 | zLinux
(Required) Specifies the host type that is defined by the operating system.

-hostport *wwpn [,...]*
(Optional) Specifies the worldwide port names (WWPNs) of a set of host ports that are mapped to the specified host. The ellipsis (...) indicates that you can specify multiple values.

-cluster *cluster_name*
(Optional) Specifies the name of the cluster that this host is assigned to.

Note: If cluster name contains a blank space, you must surround the name in quotation marks.

host_name | **-**
(Required) Specifies the name that you assigned to the host system. The name is limited to 32 byte or 16 double byte characters.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the mkhost command

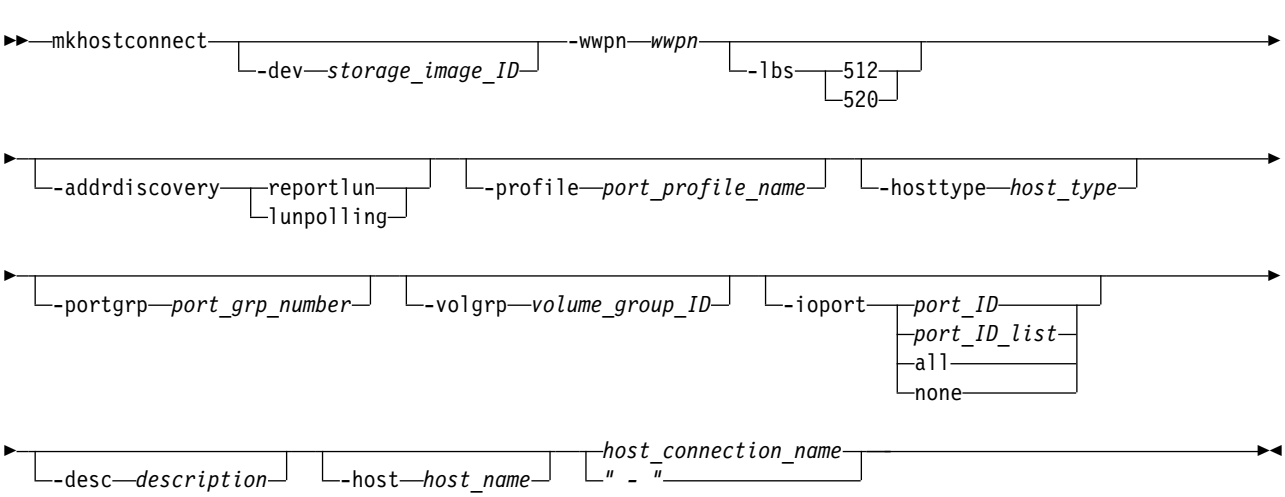
```
dscli> mkhost -type generic -hostport 210100E08B295F86 host_1
```

The resulting output

```
Aug 11 02:23:49 PST 2016 IBM DS CLI Version: 5.0.0.0 DS:IBM.2107-75FA120
Host host_1 successfully created.
```

mkhostconnect

The **mkhostconnect** command configures open systems hosts port attachments to Fibre Channel ports that are configured for SCSI-FCP topology.



Parameters

Open-systems host port attachments to Fibre Channel ports are configured for identified access mode and SCSI protocol.

Notes:

1. Ensure that you use the **-hosttype** parameter with this command to reduce input and processing time. The **-hosttype** parameter supplies the same information as if you used the following three parameters:
   - **-profile**
   - **-addrdiscovery**
   - **-lbs**
2. If you are using the HP-UX operating system, see the volume restriction that is described under the **-addrdiscovery** parameter.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified WWPN, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-wwpn *wwpn*
(Required) Specifies the worldwide port name (WWPN). The WWPN is a 16-character hexadecimal ID or a colon-separated string. The names are host attachment specific, for example, 12341234000A000F or 12:34:12:34:00:0A:00:0F.

Note: Do not create more than one host connection per WWPN, except for SUN hosts. Creating more than one host connection per WWPN (each with a different volume group) is only supported for SUN.

-lbs 512 | 520
(Optional) Specifies the logical block size that is used by the specified host system, in bytes. The block size must be compatible with the volume group type and the volume type configurations that apply to the specified host port connection. The 520-byte size is typically used by IBM i models (OS/400).

Notes:

1. Do not use the **-lbs** parameter if you specify the **-hosttype** parameter.
2. If you do not use the **-hosttype** parameter, use the **lsportprof** command to determine the block size that you need to specify.

-addrdiscovery reportlun | lunpolling
(Optional) Specifies the method for discovering logical unit number (LUN) addresses.

- The reportlun method specifies that the host system can access up to 64 000 LUNs.

  Note: Use the reportlun method only with volume groups that are designated as *mask* type. (This designation is assigned when you use the **mkvolgrp** command to create the volume group.) However, you can use the reportlun method for a *map* type, but more considerations exist if you are using an HP-UX operating system.

  For HP-UX operating systems, the number of volumes in the volume group must not exceed seven volumes. This restriction applies only when the **-addrdiscovery** parameter is set to reportlun and the associated volume group is of type scsimap256.
- The lunpolling method specifies that the host system can access up to 256 LUNs. For Sun, Linux, and Windows operating systems, the lunpolling method is typically selected.

Notes:

1. Use the lunpolling method only with volume groups that are designated as *map* type. (This designation is assigned when you use the **mkvolgrp** command to create the volume group.)
2. Do not use the **-addrdiscovery** parameter if you specify the **-hosttype** parameter.

-profile *port_profile_name*
(Optional. If you specify the **-hosttype** parameter, this parameter is not used.) Specifies the name of the host connection behavior profile. If the name includes a blank space, enclose the name with double quotation marks. For example, -profile “IBM pSeries - AIX”.

Notes:

1. Do not use the **-profile** parameter if you specify the **-hosttype** parameter.
2. If you do not use the **-hosttype** parameter, use the **lsportprof** command to obtain a list of available profiles.

-hosttype *host_type*
(Optional) Specifies information about the following three parameters:

- **-profile**
- **-addrdiscovery**
- **-lbs**

Notes:

1. When the **-hosttype** parameter is specified, do not use the **-profile**, **-addrdiscovery**, or **-lbs** parameters.
2. Use the **lshosttype** command to obtain a list of known host types.

-portgrp *port_grp_number*
(Optional) Specifies the identifier that associates two or more host ports with access to a common volume group. Port group zero is reserved for ports that are not associated with a port group.

-volgrp *volume_group_ID*
(Optional) Specifies an available volume group. This parameter accepts a fully qualified volume group ID including the storage image ID or a shortened version. The shortened version is a four-digit decimal number with no leading zeros, prefixed with the letter *V*.

A host connection uses only one volume group per storage image; that is, a single WWPN can access only one volume group per storage image.

Note: If you do not specify a volume group when a host connection is created, the value for volume group is displayed as a "-" when you enter a **lshostconnect** or **showhostconnect** command.

-ioport *port_ID port_ID_list* **|all|none**
(Optional) Specifies all, none, one, or more I/O port IDs that allow host connection access to volumes. I/O ports cannot share WWPN. Ensure that no conflicts exist with the I/O ports of existing SCSI host connections.

port_ID port_ID_list
Specifies that you can designate up to 128 ports for an open systems host attachment assignment. If you enter a list of I/O port IDs, access from the specified host connection to the specified volume group is allowed by using only the designated list of port IDs.

Note: To specify a range of port IDs, separate the port IDs with a hyphen. You must separate multiple port IDs or ranges of port IDs with a comma between each ID or range of IDs.

all Specifies that you want to add all I/O port IDs to allow the specified host connection access to the designated volume group through all the associated I/O port IDs.

none Specifies that you do not want to add any I/O ports. If you do not specify I/O ports, the storage system is configured to allow host connection access to the specified volume group. You can use any I/O port that is configured for SCSI-FCP topology. A port ID is four hexadecimal characters in the format *EEAP*, where:

EE The I/O port enclosure number in the range of 00 - 17.

A The adapter number and is specified as 1, 2, 4, or 5.

P The port number (0 - 3).

This number is prefixed with the letter *I*.

-desc *description*
(Optional) Specifies the description that you defined for the SCSI host port. The description is limited to 256 byte or 128 double-byte characters.

-host *host_name*
(Optional) The name of the host to which the host connection is assigned.

Note:

- If the host name includes a blank space, enclose the name with double quotation marks.
- You cannot use the -**host_name** parameter with the **-lbs**, **-addrdiscovery**, **-profile**, **-hosttype**, **-portgrp**, **-volgrp**, or **-ioport** parameters.

host_connection_name | **-**
(Required) Specifies the name that you assigned to the host port. The name you assign is limited to 32 byte or 16 double byte characters.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the mkhostconnect command

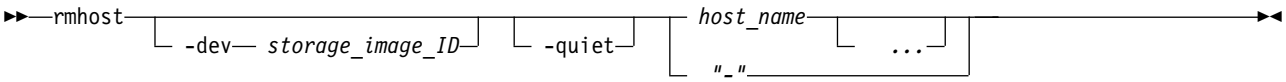
```
dscli> mkhostconnect -wwpn 12341234000A000F
-host host_1 my_host_connect
```

The resulting output

```
Aug 11 02:23:49 PST 2017 DS CLI Version: 5.0.0.0 DS:IBM.2107-75FA120
Host connection 0001 successfully created.
```

rmhost

The **rmhost** command removes a host from a storage image.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host connection ID. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-quiet
(Optional) Turns off the confirmation prompt for this command.

host_name ... | **-**
(Required) Specifies the host name to be removed. The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmhost command

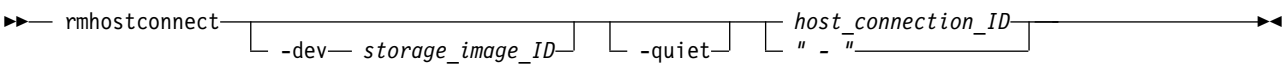
```
dscli> rmhost -quiet host_1
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2016  DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
Host host_1 successfully deleted.
```

rmhostconnect

The **rmhostconnect** command removes a SCSI host port connection from a storage image.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for all host connections, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-quiet
(Optional) Turns off the removal confirmation prompt.

host_connection_ID | **-**
(Required) Specifies the host connect ID, which is a unique identifier that uses any number from 0 - FFFE within the scope of a storage image. This parameter accepts a fully qualified ID (includes *manufacturer.machine type-serial number/hostconnectID*) or a shortened version if the **-dev** parameter is specified.

For DS8000, example of a fully qualified host connection ID: IBM.2107-75FA120/1.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Attention: Use caution when you work with connection IDs to ensure that you specified the connection that you want to delete. For example, if you intend to delete connection ID 0005 and type 000, the system deletes connection ID 0. Or, if you want to delete connection ID 0020 and type 002, the system deletes connection ID 2. The system does not consider the leading zeros, and 000 is interpreted as connection ID 0 and 002 is interpreted as connection ID 2.

Example

Invoking the rmhostconnect command

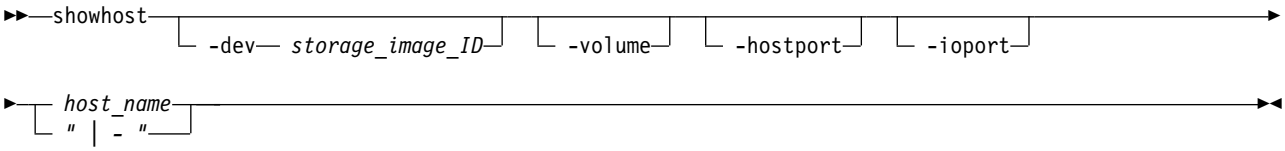
```
dscli> rmhostconnect -dev IBM.2107-75FA120 1
```

The resulting output

```
Are you sure you want to delete Host Connection IBM.2107-75FA120/1?
y/n Y
Host Connection IBM.2107-75FA120/1 successfully deleted.
```

showhost

The **showhost** command displays detailed properties of an individual host.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID that consists of manufacturer, computer type, and serial number. The storage image ID is required if you don't specify a fully qualified host connection ID. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-volume

(Optional) Displays the volume and LUN ID map for the specified host.

-hostport

(Optional) Displays the host port WWPNs and names under the specified host.

-ioport

(Optional) Displays the I/O ports for the specified host.

host_name | **-**

(Required) Displays the host information for the host names specified. Alternatively, the command accepts input from stdin when the dash (-) is specified. Fully qualified names or shortened versions are also accepted.

Example 1

Invoking the **showhost** command

```
dscli> showhost Host_1
Name                Host_1
Type                Generic
State               Online
AddrMode            scsimask
Volumes             3
Host ports          2
I/O ports           3
AddrDiscovery       reportLUN
LBS                 512
Cluster             Production01
```

Name

Displays the host name that you assigned. The name is limited to 32 byte or 16 double-byte characters.

Type

Displays the host type that is defined by the operating system. The value *Unknown* displays when the information is not available or was not specified when you created or modified the host.

State

Displays the login state of the host (online or offline). If at least one host port is in the logged-in state, the state is displayed as *online.* If no host ports were added to this host, or if all of the added host ports are logged out, the state is displayed as *offline*.

AddrMode

Displays the method that is used by the host to discover the LUNs that are available to its host ports (for example, scsimask, scsimap256, os400mask).

Volumes

Displays the number of volumes that are mapped to this host.

Host ports

Displays the number of host ports that are mapped to this host.

I/O ports

Displays the number of I/O ports this host is allowed to log in to. *All* indicates that this SCSI host is allowed to log in to any I/O ports.

AddrDiscovery

Displays the address discovery method used by this host system and host system port. One of the following values displays.

ReportLUN

Indicates that host system LUN access is limited to a maximum of 64 000 LUNs.

LUNPolling
Indicates that host system LUN access is limited to a maximum of 256 LUNs.

reportLUN16K
Indicates that host system LUN access is limited to a maximum of 16 000 LUNs.

LBS
Displays the logical block size of the volumes that the host can access. Block size is determined by the host type. Possible values include 512 and 520 (for IBM® i).

Cluster
Displays the name of cluster that the specified host is assigned to.

Example 2

Invoking the **showhost** command, with the volume parameter specified

```
dscli> showhost -volume Host_1
```

The resulting output

```
Name             Host_1
Type             Generic
State            Online
AddrMode         scsimask
Volumes          3
Host ports       2
I/O ports        3
AddrDiscovery    reportLUN
LBS              512
Cluster          Production01
================Volumes=============
ID               Lun         Access
====================================
000D             8D          Public
000E             8E          Public
000F             8F          Public
```

When you specify the **-volume** parameter, the Volume IDs, Lun, and Access parameters also display.

ID Displays the volume ID.

Lun
Displays the Lun ID that is mapped to the volume ID.

Access
Displays the accessibility status of the volume. One of the following values displays.

Public
Indicates that volume is both under the host and its cluster. This volume is shared by other hosts under the cluster.

Private
Indicates that the volume is not visible to the host's cluster.

- Indicates that if the volume is not under any cluster, the value always displays as "-".

Example 3

Invoking the **showhost** command, with the hostport parameter specified

```
dscli> showhost -hostport Host_1
```

The resulting output

```
Name             Host_1
Type             Generic
State            Online
AddrMode         scsimask
Volumes          3
Host ports       2
I/O ports        3
AddrDiscovery    reportLUN
LBS              512
Cluster          Production01
===============Host Ports===============
WWPN                   Name
=======================================
0000000000000000F      host_1_port_1
00000000000000010      host_1_port_2
```

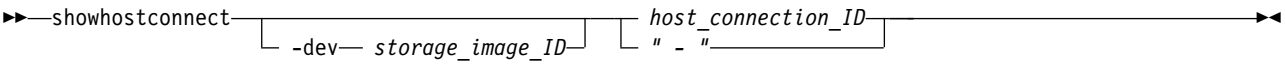
When you specify the **-hostport** parameter, the WWPN identifier and host port IDs also display.

WWPN
Displays the host port WWPN that is mapped to the host.

Name
Displays the host port name that is mapped to the host.

showhostconnect

The **showhostconnect** command displays detailed properties of a storage image host connection.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the host connection, do not set the *devid* variable in your profile or through the **setenv** command. The storage image ID is also required if the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

host_connection_ID | **-**
(Required) Specifies a fully qualified host connection ID, which includes the manufacturer, machine type, and serial number if the **-dev** parameter is not used. The host connection *ID* is a unique identifier (0 - FFFE) within the scope of a storage image.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **showhostconnect** command.

Invoking the showhostconnect command

```
dscli> showhostconnect -dev IBM.2107-75FA120 1
```

The resulting output

Name	ID	WWPN	HostType	LBS	Addrdis covery
My host port 1	IBM.2107-75FA120/1	3007ACF3 0A2399E0	Unknown	512	reportLUN

Profile	Portgrp	VolgrpID	Atchtopo	ESSIOport	Speed	Desc	Host
IBM pSeries - AIX	0	100	-	I0111,I0121 I0211,I0221	Unknown	SCSI1	Production 01

Report field definitions

Name
Indicates the host connection SCSI port nickname.

The name is limited to 32 byte or 16 double-byte characters.

ID Indicates a fully qualified host connection ID.

The value that is represented by the *host_connection_ID* parameter is a unique identifier (0 - FFFE) within the scope of a storage system.

WWPN
Indicates the worldwide port name (WWPN) for the designated host system port.

HostType
Indicates the name of the host type.

"Unknown" is displayed when the information is not available, and indicates that the host type was not specified when the host connection was created or modified.

LBS
Indicates the logical block size that is used by this host system and the host system port.

The logical block setting must be compatible with the volume group type that is configured for volume access. The 520 block size is typically used for IBM System i® host system attachments.

Addrdiscovery
Indicates the LUN address discovery method that is used by this host system and the host system port.

The LUN Discovery method must be compatible with the volume group type that is configured for volume access.

The Poll LUNs method enables access to a maximum of 256 LUNs. The Report LUNs method enables access to a maximum of 64 000 LUNs.

Profile
Indicates the name of the host connection behavior profile.

Portgrp
Indicates the host port group ID. The ID ties together a group of SCSI host port objects that are accessing a common volume group. If the port group value is set to zero, the host port is not associated with any port group.

VolgrpID
Indicates the volume group ID. This ID is a unique identifier within the DS8000 for the SCSI volume group that the specified SCSI host port is assigned to.

Atchtopo
The atchtopo is an attribute of the I/O port and, under certain conditions, might display a value that is inconsistent with the value reported by the **lsioport** command. Because this inconsistency cannot

be corrected in all cases, the atchtopo attribute always displays as a " - " value. To obtain the I/O port topology, use the **lsioport** or **showioport** commands.

ESSIOport
Indicates the array of port IDs that the designated host port is logged in to.

A port ID is prefixed with the letter "I" and consists of four hexadecimal characters in the format *EEAP*, where:

- *EE* is an I/O port enclosure number in the range of 00 - 17.
- *A* is the adapter number and is specified as 1, 2, 4, or 5.
- *P* is the port number (0 - 3).

Speed
Speed is an attribute of the I/O port and, under certain conditions, might display a value that is inconsistent with the value reported by the **lsioport** command. Because this inconsistency cannot be corrected in all cases, the speed attribute always displays a value of "Unknown". To query the speed of an I/O port, use the **lsioport** or **showioport** commands.

Desc
Indicates the description that you defined for the SCSI host port. The description is limited to 256 byte or 128 double-byte characters.

Host
Indicates the host name that is used in the Storage Management GUI, which is different from the nickname that is displayed in the **Name** column. If a name is not available, a value of "-" displays.

Host port commands

Use the host port commands to configure, modify, or delete, or display open system host ports.

The following host port commands are available.

chhostport
Modifies an open systems host port.

mkhostport
Configures an open systems host port.

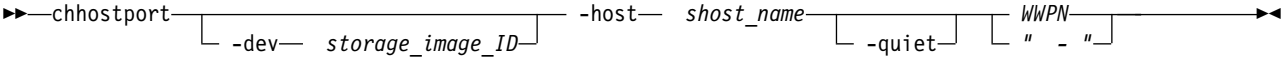
lshostport
Displays a list of host ports for a specified storage image and status information for each host port in the list.

rmhostport
Deletes an open systems host port.

showhostport
Displays detailed properties of an individual host port.

chhostport

The **chhostport** command modifies an open systems host port.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host port. It is also

required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-host *host_name*
(Required) Specifies the name of the host that the host port is reassigned to. If you specify a host name with a space, you must add double quotation marks around the host name.

-quiet
(Optional) Turns off the confirmation prompt for this command.

WWPN | **-**
(Required) Specifies the worldwide port name (WWPN) of the SCSI host port on the host system that is connected to the storage system, either directly or through a switch.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the chhostport command

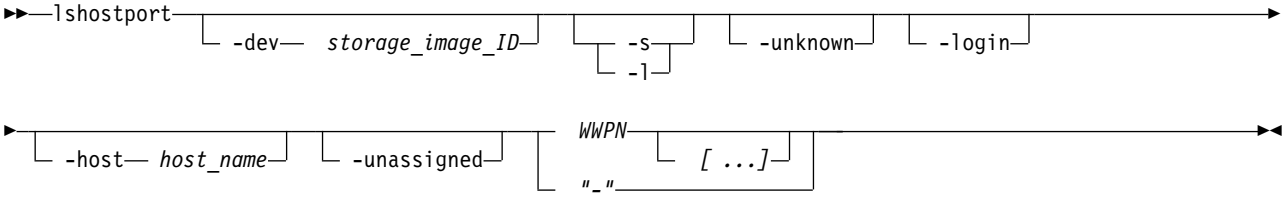
```
dscli> chhostport -host host_2 12341234000A000F
```

The resulting output

```
Aug 11 02:23:49 PST 2018 DS CLI Version: 5.0.0.0 DS:IBM.2107-75FA120
Host Port 12341234000A000F successfully modified.
```

lshostport

The **lshostport** command displays a list of host ports for a specified storage image and status information for each host port in the list.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host port. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-s
(Optional) Displays the worldwide port name (WWPN) of the host ports. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output and user-defined description for each host port in the list. You cannot use the **-l** and the **-s** parameters together.

-unknown
(Optional) Displays a list of logged in host ports (WWPNs) that are not recognized as being associated with the designated storage system.

Note: You cannot use the **-unknown** parameter with the **-login**, **-host**, **-unassigned,** or **WWPN** parameters.

-login
(Optional) Displays a list of host ports (WWPNs) that are logged in and sorted by the ESS I/O port IDs for known connections. The report displays one line of information per connection. However, no information is displayed for a FICON connection.

Notes:

1. Known logins are logins that you created for host connections, Remote Mirror and Copy paths, and anonymous connections.
2. You cannot use the **-login** parameter with the **-unknown**, **-host**, **-unassigned,** or **WWPN** parameters.

-host *host_name*
(Optional) Displays host port information that is associated with the specified host.

Notes:

1. If you specify a host name with a blank space, enclose the name with double quotation marks.
2. The **-host** parameter is not compatible with the **-unassigned** parameter.

-unassigned
(Optional) Displays host port information that is not associated with other hosts.

Note: The **-unassigned** parameter is not compatible with the **-host, -unknown,** or **-login** parameters.

WWPN [...] | **-**
(Optional) Displays host port information for one or more specified host ports (WWPNs). If you specify multiple WWPNs, separate each one with a blank space. You can specify the fully qualified WWPN or the shortened version of the WWPN.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 1

For this command, the results are shown in table format to provide clarity. The actual reports do not display as tables.

Invoking the **lshostport** command, without the **-login** and **-unknown** parameters specified

```
dscli> lshostport -l
```

The resulting output

WWPN	Name	Type	State	Volumes	AddrMode	I/O ports	Host
210100E08B 295F86	Name_1	VMWare	logged in	8	scsimap256	32 (all)	host_1
310106E04 B2D5E45	-	VMWare	logged out	6	scsimap256	3	-
310106E04B 2D5F10	-	-	-	0	-	0	-

Report field definitions

WWPN*
Displays the WWPN for the specified host port.

Name
Displays the name of the host port.

Type
Displays the host type that is defined by the operating system and platform of the host.

When information is not available, the output displays "Unknown".

State
Displays the login state of the host. One of the following values displays.

Logged in
Indicates that a host port is created for an available physical connection.

Logged out
Indicates that a host port is created but a physical connection is not established and the connection is not available.

unknown
Indicates that a host port is created but the state is invalid (not logged in or logged out).

"-"
Indicates that a physical connection is established but a host port is not created.

Volumes+
Displays the number of volumes that are mapped to this host port.

AddrMode+
Displays the method that is used by the host to discover the LUNs that are available to its host ports. One of the following values display.

scsimask
Indicates a SCSI mask volume group. This option is available if the host adapter supports 4-byte LUN addresses.

scsimap256
Indicates a SCSI-MAP 256 volume group.

os400mask
Indicates an OS/400 mask volume group. The IBM i host system typically uses fixed block volumes of 520-byte logical block size. This option is available only if the host adapter supports 4-byte LUN addresses.

Note: This volume group is also referred to as SCSI520-MASK. When an error message is displayed for the os400mask, SCSI520-MASK is referenced instead.

I/O ports+
Displays the number of I/O ports that this host is allowed to log into. The value of "all" indicates that this SCSI host is allowed to log into any I/O port.

Host+
Displays the name of the host that the specified host port is assigned to. The value of "-" indicates that the host port is not assigned to any host.

Key:

* Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

Example 2

Invoking the **lshostport** command with the **-unknown** parameter specified

```
dscli> lshostport -dev IBM.2107-75FA120 -unknown
```

The resulting output

WWNN	WWPN	ESSIOport
5005076306FFD341	5005076306FFD341	I0221

Report field definitions

When you specify the **-unknown** parameter, the WWPN and WWPN identifiers and ESSIOport IDs also display.

WWNN
Displays the WWNN for the specified host system.

WWPN
Displays the WWPN for the specified host port.

ESSIOport
Displays the port ID to which the designated host port is logged in to. A port ID is prefixed with the letter I and consists of four hexadecimal characters in the format *EEAP*, where:

EE An I/O port enclosure number in the range of 00 - 17.

A The adapter number and is specified as 1, 2, 4, or 5.

P The port number (0 - 3).

Example 3

Invoking the **lshostport** command with the **-login** parameter specified

```
dscli> lshostport IBM.2107-75FA120 -login
```

The resulting output

WWNN	WWPN	Host	I/O ports	ESSIOport
5005076306FFD341	5005076306FFD341	-	32 (all)	I0221

When you specify the **-login** parameter, the WWPN and WWPN identifiers, host, I/O ports and ESSIOport IDs also display.

Report field definitions

WWNN
Displays the WWNN for the specified host system.

WWPN
Displays the WWPN for the specified host port.

Host
Displays the the name of the host that the specified host port is assigned to. A hyphen (-) displays if the host port is not associated with a host.

I/O ports
Displays the number of I/O ports that this host is allowed to log into. *All* displays if the specified SCSI host is allowed to log into any I/O port.

ESSIOport
Displays the port ID to which the designated host port is logged in to. A port ID is prefixed with the letter I and consists of four hexadecimal characters in the format *EEAP*, where:

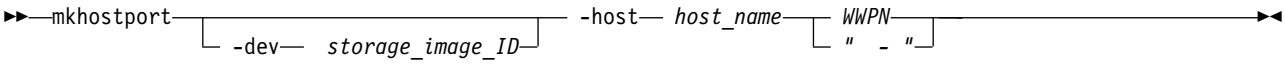
EE An I/O port enclosure number in the range of 00 - 17.

A The adapter number and is specified as 1, 2, 4, or 5.

P The port number (0 - 3).

mkhostport

The **mkhostport** command configures an open systems host port.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host port. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-host *host_name*
(Required) Specifies the name of the host. If you specify a host name with a blank space, enclose the name with double quotation marks.

WWPN | **-**
(Required) Specifies the worldwide port name (WWPN) of the SCSI host port on the host system that is connected to the storage system, either directly or through a switch.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the mkhostport command

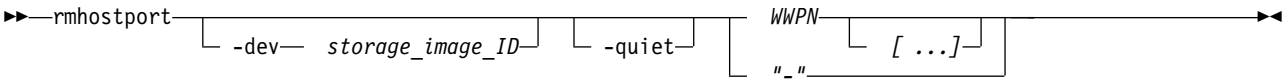
```
dscli> mkhostport -host host_1 12341234000A000F
```

The resulting output

```
Aug 11 02:23:49 PST 2016 DS CLI Version: 5.0.0.0 DS:IBM.2107-75FA120
Host Port 12341234000A000F successfully created.
```

rmhostport

The **rmhostport** command deletes an open systems host port.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host port. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-quiet
(Optional) Turns off the confirmation prompt for this command.

WWPN [...] | **-**
(Required) Specifies the worldwide port name (WWPN) of the SCSI host port on the host system that is connected to the storage system, either directly or through a switch. The ellipsis (...) indicates that you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmhostport command

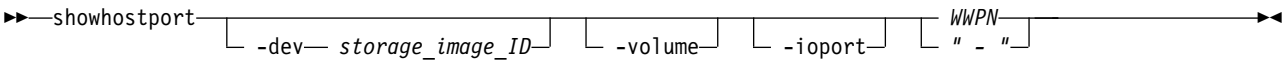
```
dscli> rmhostport -quiet 12341234000A000F
```

The resulting output

```
Aug 11 02:23:49 PST 2018 DS CLI Version: 5.0.0.0 DS:IBM.2107-75FA120
Host Port 12341234000A000F successfully deleted.
```

showhostport

The **showhostport** command displays detailed properties of an individual host port.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified host port. It is also required if you did not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-volume
(Optional) Displays volume and LUN ID map for the specified host port.

-ioport
(Optional) Displays I/O ports for the specified host port.

WWPN | **-**
(Required) Displays host port information for the specified host port WWPNs. Either the fully qualified WWPN or shortened name is acceptable. This parameter accepts input from stdin when the dash (-) is specified.

Example 1

Invoking the **showhostport** command.

```
dscli> showhostport 210100E08B295F86
```

The resulting output

```
WWPN              210100E08B295F86
Name              Host_Port_1
Type              generic
State             logged in
Volumes           3
AddrMode          scsimask
```

```
I/O ports          3
AddrDiscovery      reportLUN
LBS                512
Host               Host_1
```

Report field definitions

WWPN
Displays the WWPN of the host port.

Name Displays the unique ID that is assigned by the user to the host port.

Type Displays the host type as defined by the operating system and platform. If the information is not available, the value always displays as `unknown`.

State Displays the login state of the host.

logged in
Indicates that a host port is created for an available physical connection.

logged out
Indicates that a host port is created, but no physical connection is established or the connection is not available.

unknown
Indicates that a host port is created, but the state is invalid (neither logged in or logged out).

- A physical connection is established, but the host port is not created.

Volumes
Displays the number of volumes that are mapped to the specified host.

AddrMode
Displays the method that is used by the host to discover the LUNs that are available to its host ports; for example, `scsimask`, `scsimap256`, or `os400mask`.

I/O ports
Displays the number of I/O ports that the specified host is allowed to log into (`all` indicates that the specified host is allowed to log into any I/O port).

AddrDiscovery
Displays the LUN address discovery method used by the host system and host system port; for example, `reportLUN`, `LUNPolling`, or `reportLUN16K`.

LBS Displays the logical block size of the volumes that the host can access. Block size is determined by the host type. For IBM i, the possible values include `512` and `520`.

Host Displays the name of the host that the specified host port is assigned. If the host port is not associated with a host, the value always displays as "-".

Example 2

Invoking the **showhostport** command with the **-volume** parameter specified

```
dscli> showhostport -volume 210100E08B295F86
```

The resulting output

```
WWPN             210100E08B295F86
Name             Host_Port_1
Type             generic
State            logged
Volumes          3
AddrMode         scsimask
I/O ports        3
```

```
AddrDiscovery    reportLUN
LBS              512
Host             Host_1
===============Volumes==================
ID            Lun       Access
=========================================
000D          8D        Public
000E          8E        Public
000F          8F        Public
```

When you specify the **-volume** parameter, the Volume IDs, Lun, and Access parameters also display.

ID Displays the volume ID.

Lun
Displays the Lun ID that is mapped to the volume ID.

Access
Displays the accessibility status of the volume. One of the following values displays.

Public
Indicates that volume is both under the host and its cluster. This volume is shared by other hosts under the cluster.

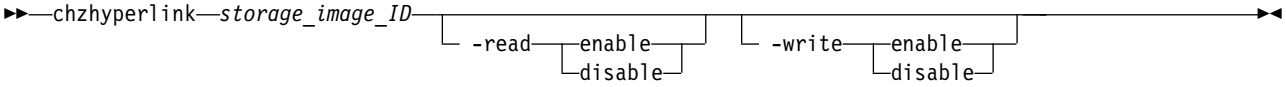
Private
Indicates that the volume is not visible to the host's cluster.

- Indicates that if the port's host is not under any cluster, the value always displays as "-".

chzhyperlink

The **chzhyperlink** command modifies the zHyperLink control switch for the specified storage image.

Note: zHyperLink connections are not available on DS8882F systems.



Parameters

storage_image_ID
(Required) Specifies the storage image ID, which consists of the values for manufacturer, machine type, and serial number.

Note: If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

-read enable | disable
(Optional) Specifies the setting of the read control switch.

-write enable | disable
(Optional) Specifies the setting of the write control switch.

Note: By default, the zHyperLink facility is disabled in z/OS. To enable z/OS for zHyperLink read or write processing, use the SETIOS MVS™ system command after an IPL, or use the IECIOSxx PARMLIB to enable zHyperLink processing during IPL.

SETIOS MVS Command
To enable zHyperLink read and write processing on the IBM z host system after an IPL, enter the following MVS commands:

```
SETIOS ZHPF=YES
SETIOS ZHYPERLINK,OPER={READ|WRITE|ALL}
```

For more information about the SETIOS MVS command, refer to the z/OS MVS System Commands.

IECIOSxx parmlib

To enable zHyperLink read and write processing during an IPL, include the following statements in the IECIOSxx PARMLIB member on the IBM z host system:

```
ZHPF=YES
ZHYPERLINK,OPER={READ|WRITE|ALL}
```

For more information about the IECIOSxx PARMLIB member, refer to the z/OS MVS Initialization and Tuning Reference.

Example

Invoking the chzhyperlink command

```
dscli> chzhyperlink -read enable IBM.2107-75FA120
```

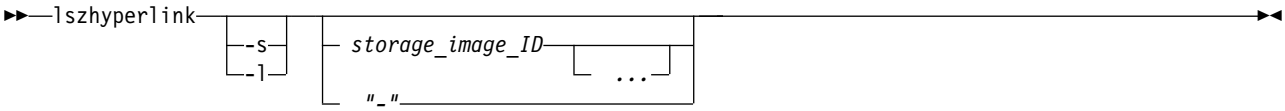
Output

```
CMUC00519I chzhyperlink: zHyperLink read is successfully modified.
```

lszhyperlink

The **lszhyperlink** command displays the status of zHyperLink control switches for specified storage images.

Note: zHyperLink connections are not available on DS8882F systems.



Parameters

-s

(Optional) Displays only the storage image IDs. You cannot use the -l and the -s parameters together.

-l

(Optional) Displays the default output for each storage image ID in the list. You cannot use the -l and the -s parameters together.

storage_image_ID ... | -

(Required) Specifies one or more storage image IDs, which consists of the values for manufacturer, machine type, and serial number. You must separate multiple IDs with a space between each ID.

Note: If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode. You cannot specify ID ranges.

Example

Invoking the lszhyperlink command

```
dscli> lszhyperlink
```

The resulting output

```
ID                  Read     Write
==================================
IBM.2107-75FA120    disable  enable
```

Report field definitions

ID Indicates the storage image ID that consists of the manufacture, machine type, and serial number.

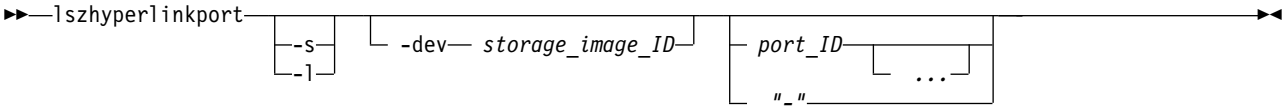
Read
Indicates the state of the zHyperLink read control switch.

Write
Indicates the state of the zHyperLink write control switch.

lszhyperlinkport

The **lszhyperlinkport** command displays a list of zHyperLink ports for the specified storage image.

Note: zHyperLink connections are not available on DS8882F systems.



Parameters

-s
(Optional) Displays only the port IDs. You cannot use the -l and the -s parameters together.

-l
(Optional) Displays the default output for each port ID in this list. You cannot use the -l and the -s parameters together.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified port ID. It is also required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

port_ID ... | **-**
(Optional) Displays the ports that match the specified IDs. A port ID is prefixed with the letter "HL" and consists of four hexadecimal characters.

Example

Invoking the lszhyperlinkport command

```
dscli> lszhyperlinkport
```

The resulting output

```
ID           State                     Loc                  Speed
=================================================================
HL0028       Connected        U1500.1B3.RJBAY03-P1-C7-T3       8 GB/s
HL0029       Connected        U1500.1B3.RJBAY03-P1-C7-T4       8 GB/s
HL0038       Disconnected     U1500.1B4.RJBAY04-P1-C7-T3       8 GB/s
HL0039       Disconnected     U1500.1B4.RJBAY04-P1-C7-T4       8 GB/s
```

Report field definitions

ID Indicates the zHyperLink port ID.

State
Indicates the state of the zHyperLink port in the I/O enclosure PCIe/PCN card. One of the following values display: Connected, Disconnected, Offline, or Fenced.

Loc
Indicates the location of the zHyperLink port in the I/O enclosure PCIe/PCN card.

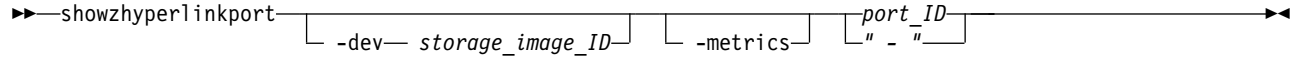
Speed
Indicates the speed of the link of the zHyperLink port in the I/O enclosure PCIe/PCN card.

Width
Indicates the number of PCIe lanes operational on the zHyperLink port in the I/O enclosure PCIe/PCN card.

showzhyperlinkport

The **showzhyperlinkport** command displays the detailed properties of an individual zHyperLink port.

Note: zHyperLink connections are not available on DS8882F systems.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified port ID. It is also required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-metrics
(Optional) Displays the port information and performance metrics for the specified port.

port_ID | **-**
(Required) Display port information for the specified port.

Example 1

Invoking the showzhyperlinkport command

```
dscli> showzhyperlinkport -dev IBM.2107-75DMC81 HL0038
```

The resulting output

```
ID           HL0038
State        Connected
Loc          U1500.1B8.RJBAY08-P1-C7-T9
Speed        8 GB/s
```

Report field definitions

ID Indicates zHyperLink port ID.

State
Indicates the state of the zHyperLink port in the I/O enclosure PCIe/PCN card. One of the following values display: Connected, Disconnected, Offline, or Fenced.

Loc
Indicates the location of the zHyperLink port in the I/O enclosure PCIe/PCN card.

Speed
Indicates the speed of the link of the zHyperLink port in the I/O enclosure PCIe/PCN card.

Example 2

Invoking the showzhyperlinkport command to show port performance metrics

```
dscli> showzhyperlinkport -dev IBM.2107-75DMC81 -metrics HL0028
```

The resulting output

```
ID                  HL0028
Date                Fri Jun 23 11:26:15 PDT 2017
TxLayerErr          2
DataLayerErr        3
PhyLayerErr         4
============================================
Lane         RxPower(dBm)    TxPower(dBm)
============================================
1             0.1845           -0.2909
2              -0.41           -0.0682
3              0.114           -0.4272
```

Report field definitions (with the `-metrics` parameter)

ID Indicates the zHyperLink port ID.

Date
Indicates the timestamp when the performance metrics were captured.

TxLayerErr
Indicates the number of bad transaction layer packets.

DataLayerErr
Indicates the number of bad data layer packets.

PhyLayerErr
Indicates the number of bad physical layer packets.

Lane
Indicates the number of the lane.

RxPower (dBm)
Indicates the link transmit optical power in dBm.

TxPower(dBm)
Indicates the link receiver optical power in dBm.

Network port commands

Network port commands are available for configuring and displaying network ports information for your storage system.

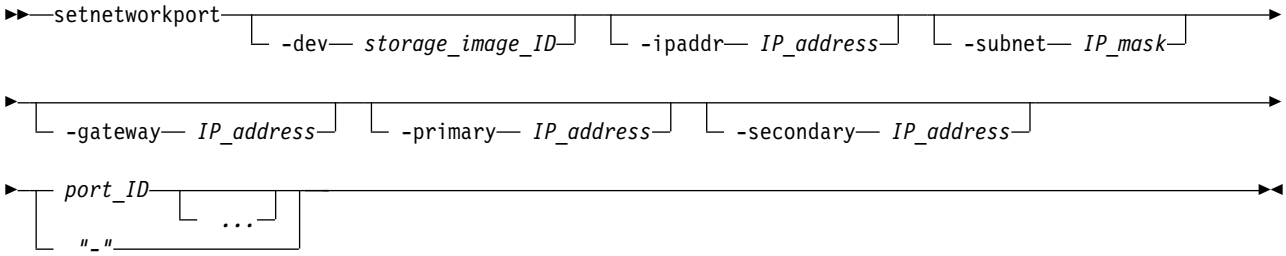
lsnetworkport
Displays a report that lists all network ports that are installed in the specified storage system.

setnetworkport
Configures one network port for TCP/IP connections.

shownetworkport
displays detailed properties of an individual network port ID.

setnetworkport

The **setnetworkport** command configures one network port for TCP/IP connections.



Parameters

Notes:

1. One or more of the optional parameters, **-ipaddr**, **-subnet**, **-gateway**, **-primary**, or **-secondary**, must be specified.
2. Only one port can be set or modified at a time.

-dev *storage_image_ID*
(Optional) Accepts a fully qualified storage image ID. The storage image ID consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified port ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter will temporarily override any defined value for *devid* for the current command.

-ipaddr *IP_address*
(Optional) Specifies the TCP/IP address for the designated port.

Note: Ports that are on the same server must be on different subnets.

-subnet *IP_mask*
(Optional) Specifies the TCP/IP subnet mask for the designated port.

Note: Ports that are on the same server must be on different subnets.

-gateway *IP_address*
(Optional) Specifies the gateway TCP/IP address for the designated port.

-primary *IP_address*
(Optional) Specifies the primary DNS TCP/IP address for the designated port.

-secondary *IP_address*
(Optional) Specifies the secondary DNS TCP/IP address for the designated port.

port_ID ... | **-**
(Required) Specifies the network port ID. This parameter accepts a fully qualified port ID, which includes the storage image ID, or a shortened version without the storage image ID when the **-dev** parameter is specified.

A port ID is prefixed with letter *I* and consists of four hexadecimal characters with the binary bit definition of *trre eeee aaaa pppp* (value is separated for readability), where:

- *t* = port type (1= network port)
- *rr* = reserved bits and set to zero
- *eeeee* = enclosure number (value specified is 18 - 1F)
- *aaaa* = adapter number

- *pppp* = port number

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the setnetworkport command

This example shows how you would configure network port (IBM.2107-1300321/I9801) to use IP address 192.168.5.5.

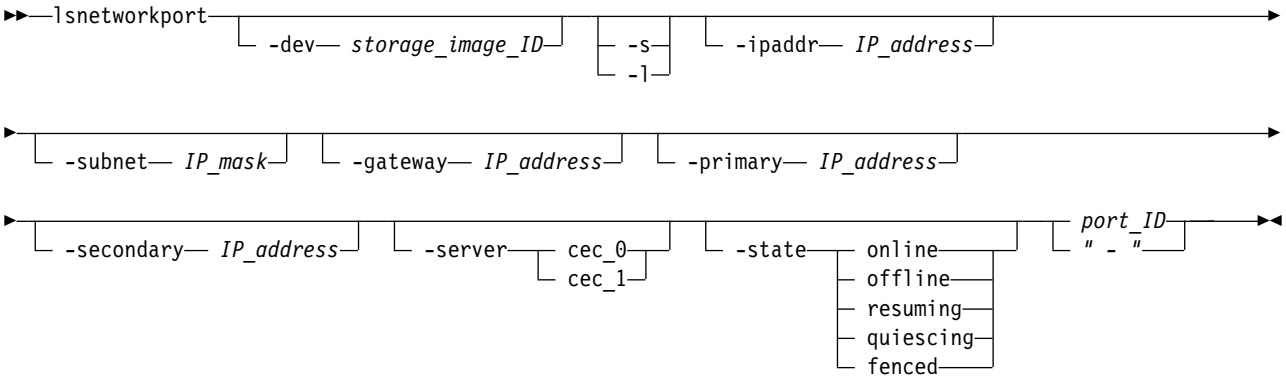
```
dscli> setnetworkport -dev IBM.2107-1300321 -ipaddr 192.168.5.5 I9801
```

The resulting output

```
CMUC00250I setnetworkport: Network port I9801 successfully configured.
```

lsnetworkport

The **lsnetworkport** command displays a report that lists all network ports that are installed in the specified storage system.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified port ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-s
(Optional) Specifies that only the port IDs be displayed for the designated storage system. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output and the values for the owning server, port speed, and port location. You cannot use the **-l** and the **-s** parameters together.

-ipaddr *IP_address*
(Optional) Specifies the TCP/IP address for the designated port.

Note: Ports that are on the same server must be on different subnets.

-subnet *IP_mask*
(Optional) Specifies the TCP/IP subnet mask for the designated port.

Note: Ports that are on the same server must be on different subnets.

-gateway *IP_address*
(Optional) Specifies the gateway TCP/IP address for the designated port.

-primary *IP_address*
(Optional) Specifies the primary DNS TCP/IP address for the designated port.

-secondary *IP_address*
(Optional) Specifies the secondary DNS TCP/IP address for the designated port.

-server cec_0 | cec_1
(Optional) Displays the server that is associated with the I/O port. One of the following values is displayed:

cec_0 Specifies that the user port is connected to processor node 0 and that the I/O port is associated with server 0.

cec_1 Specifies that the user port is connected to processor node 1 and that the I/O port is associated with server 1.

-state online | offline | resuming | quiescing | fenced
(Optional) Displays network ports that are in the specified state. One of the following values is displayed:

online Specifies that the network port is enabled.

offline
Specifies that the network port is disabled.

resuming
Specifies that the network port is in the process of becoming enabled.

quiescing
Specifies that the network port is in the process of becoming disabled.

fenced
Specifies that the network port is configured but there is a problem and the port cannot be used.

port_ID | **-**
(Required) Displays I/O ports that match the specified IDs. This parameter accepts a fully qualified port ID, which includes the storage image ID, or a shortened version without the storage image ID when the **-dev** parameter is specified.

A port ID is prefixed with letter *I* and consists of four hexadecimal characters with the binary bit definition of *trre eeee aaaa pppp* (value is separated for readability), where:

- *t* = port type (1= network port)
- *rr* = reserved bits and set to zero
- *eeeee* = enclosure number (value specified is 18 - 1F)
- *aaaa* = adapter number
- *pppp* = port number

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lsnetworkport** command using the **-l** parameter.

Invoking the lsnetworkport command

```
dscli> lsnetworkport
-dev IBM.2107-75FA120 -l
```

The resulting output

ID	IP Address	Subnet Mask	Gateway	Primary DNS
IBM.2107-75FA120/I9801	192.168.5.5	255.255.255.0	192.168.5.5	192.168.5.5

Secondary DNS	State	Server	Speed	Type	Location
192.168.5.5	Online	00	1 Gb/sec	Ethernet-Copper	U7879.001.03AA6D0-P0-C0-T0

Report field definitions

ID Specifies the port ID.

IP Address
Specifies the TCP/IP address for the designated port.

Subnet Mask
Specifies the TCP/IP subnet mask for the designated port.

Gateway
Specifies the gateway TCP/IP address for the designated port.

Primary DNS
Specifies the primary DNS TCP/IP address for the designated port.

Secondary DNS
Specifies the secondary DNS TCP/IP address for the designated port.

State Specifies the state of the network ports. One of the following values can be displayed in this field:

online Specifies that the network port is enabled.

offline
Specifies that the network port is disabled.

resuming
Specifies that the network port is in the process of becoming enabled.

quiescing
Specifies that the network port is in the process of becoming disabled.

fenced
Specifies that the network port is configured but there is a problem and the port cannot be used.

Server Specifies which server that the I/O port is attached to.

Speed Specifies the network port interface speed.

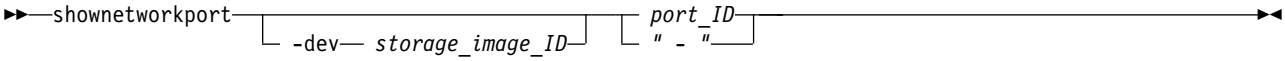
Type Specifies the network port interface type.

Location
Specifies the network port location using the format: *Utttt.mmm.ppssssss-Pn-Cn-Tn*

- *Utttt.mmm.ppssssss* identifies the location of the processor node.
- *Pn* identifies the planer number.
- *Cn* identifies the card number.
- *Tn* identifies the port number.

shownetworkport

The **shownetworkport** command displays detailed properties of an individual network port ID.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified port ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

port_ID | **-**
(Required) Displays I/O ports that match the specified IDs. This parameter accepts a fully qualified port ID, which includes the storage image ID, or a shortened version without the storage image ID when the **-dev** parameter is specified.

A port ID is prefixed with letter *I* and consists of four hexadecimal characters with the binary bit definition of *trre eeee aaaa pppp* (value is separated for readability), where:

- *t* = port type (1= network port)
- *rr* = reserved bits and set to zero
- *eeeee* = enclosure number (value specified is 18 - 1F)
- *aaaa* = adapter number
- *pppp* = port number

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **shownetworkport** command using the **-fullid** command parameter.

Invoking the shownetworkport command

```
dscli> shownetworkport -fullid -dev IBM.2107-1300321 I9801
```

The resulting output

ID	IP Address	Subnet Mask	Gateway	Primary DNS
IBM.2107-1300321 /I9801	192.168.5.5	255.255.255.5	192.168.5.5	192.168.5.55

Secondary DNS	State	Server	Speed	Type	Location
192.168.5.55	Online	00	1 Gb/sec	Ethernet-Copper	U7879.001. 03AA6D0-P0-C0-T0

Report field definitions

ID Specifies the port ID.

IP Address
Specifies the TCP/IP address for the designated port.

Subnet Mask
Specifies the TCP/IP subnet mask for the designated port.

Gateway
Specifies the gateway TCP/IP address for the designated port.

Primary DNS
Specifies the primary DNS TCP/IP address for the designated port.

Secondary DNS
Specifies the secondary DNS TCP/IP address for the designated port.

State Specifies the state of the network ports. One of the following values can be displayed in this field:

online Specifies that the network port is enabled.

offline
Specifies that the network port is disabled.

resuming
Specifies that the network port is in the process of becoming enabled.

quiescing
Specifies that the network port is in the process of becoming disabled.

fenced
Specifies that the network port is configured but there is a problem and the port cannot be used.

Server Specifies which server that the network port is attached to.

Speed Specifies the network port interface speed.

Type Specifies the network port interface type.

Location
Specifies the network port location in the format *Utttt.mmm.ppssssss-Pn-Cn-Tn* where

- *Utttt.mmm.ppssssss* identifies the location of the processor node.
- *Pn* identifies the planer number.

- *Cn* identifies the card number.
- *Tn* identifies the port number.

Offload file commands

Offload file commands are used to provide a report that contains information about who logged in, when they logged in and what the user did during their session. Also, the offload commands are used for offloading specified data files.

The following offload file commands are available:

offloadauditlog
Generates an activity report for a console that includes basic information, such as, a list of who logged in, when they logged in, and what they did during their session.

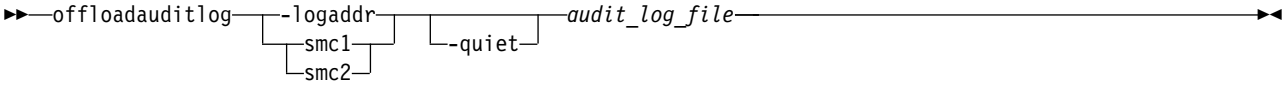
offloadfile
Offloads the specified set of data files.

offloadauditlog

The **offloadauditlog** command provides an activity report for a console (identified as smc1 or smc2).

The report includes basic information, such as, a list of who logged in, when they logged in, and what they did during their session. In addition, a log of service actions (phone connection started, phone connection ended management console session started, management console session ended) is appended to the end of the audit log.

Note: The **offloadauditlog** command cannot be run from the Embedded DS CLI window.



Parameters

Notes:

1. Only users with administrator or secadmin authority are authorized to use this command.
2. A separate log entry is added each time a resource is created, deleted, or modified. Entries are added to the audit file only after the operation has completed.
3. You must periodically extract the log using the **offloadauditlog** command and save the log file in a directory of your choice. The log file is automatically reduced (old entries removed first) by the subsystem so that it does not consume more than 50 megabytes of disk storage.

   When the log is 60% full, an entry ("Audit_Log_At_60%") is placed in the audit log. Another entry is added when the log is 75% ("Audit_Log_At_75%") full. At 100%, the log is reduced to 50% full.
4. In the service actions section of the report, there might be cases where a management console session started entry might not have a corresponding management console session ended entry.
5. As noted in the report example, the service action report begins after the line of text that states:

   ```
   "-----BEGIN SERVICE AUDIT LOG-----"
   ```

-logaddr *smc1*|*smc2*
(Required) Specifies that the audit log be offloaded for the designated storage management console. The designated storage management console must be configured and available to offload the audit log successfully.

-quiet
(Optional) Turns off the confirmation prompt for replacing an existing audit log.

audit_log_file
(Required) Specifies the file name to which the audit log entries are extracted.

Note: If you specify a file name that contains prior log entries, these entries are overwritten with the current data.

Example

Note: The following example displays only a portion of an actual Service action report appended to the end of the Audit log report.

Invoking the `offloadauditlog` command

```
dscli> offloadauditlog -logaddr smc1 auditlog-200509.txt
```

The resulting output

```
Audit log successfully offloaded from smc1 to file auditlog-200509.txt.
```

Representative report

The following lines are an example of the report information that is extracted when you use the **offloadauditlog** command (the wrapping is done simply for clarity and is not representative of your actual report):

```
U,2005/10/04 15:08:46:834 MST,admin,1,,W,1002,User_Login_Fail,,1,
"IP = an_ip_address"
U,2005/10/04 15:29:37:432 MST,admin,1,,W,1001,User_Login_Expire,,0,
"IP = an_ip_address"
U,2005/10/04 15:32:56:979 MST,admin,1,,N,1000,User_Login,,0,
"IP = an_ip_address"
U,2005/10/04 15:34:21:020 MST,admin,1,,N,1000,User_Login,,0,
"IP = an_ip_address"
U,2005/10/05 16:54:32:171 MST,admin,1,,N,1103,
User_Password_Change,,be741104,"userName = admin"
S,2005/10/06 00:01:10:239 MST,,1,,W,1200,Audit_Log_At_60%,,,""
U,2005/10/06 00:23:09:817 MST,admin,1,IBM.2107-AZ12341,N,2050,
Array_Create,A0,0,"A0"
U,2005/10/06 00:23:10:518 MST,admin,1,IBM.2107-AZ12341,N,2060,
Rank_Create,R0,-1,"R0"
U,2005/10/06 00:23:12:110 MST,admin,1,IBM.2107-AZ12341,N,2070,
XPool_Create,P0,0,"P0"
U,2005/10/06 00:23:12:761 MST,admin,1,,N,2073,XPool_Assign_Rank,,,""
U,2005/10/06 00:23:16:947 MST,admin,1,IBM.2107-AZ12341,N,2090,
Volume_Create,1000,0,"1000"
U,2005/10/06 00:23:17:187 MST,admin,1,IBM.2107-AZ12341,N,2090,
Volume_Create,1001,,"1001"
S,2005/10/06 00:23:24:508 MST,,1,,W,1201,Audit_Log_At_75%,,,""
U,2005/10/06 12:47:16:345 MST,admin,1,IBM.2107-AZ12341,N,2092,
Volume_Delete,2005,0,"2005"
U,2005/10/06 12:47:16:656 MST,admin,1,IBM.2107-AZ12341,N,2092,
Volume_Delete,2006,-1,"2006"

-----BEGIN SERVICE AUDIT LOG-----

U,2007/08/21 12:05:24:000 MST,CE,1,IBM.2107-75R0830,N,8020,
Web_SM_session_start,Web_SM_session_started,,,
U,2007/08/21 12:05:27:000 MST,CE,1,IBM.2107-75R0830,N,8022,
Web_SM_session_ended,,,,
U,2007/08/24 11:59:36:000 MST,CE,1,IBM.2107-75R0830,N,8020,
Web_SM_session_start,Web_SM_session_started,,,
U,2007/08/24 12:02:31:000 MST,CE,1,IBM.2107-75R0830,N,8022,
Web_SM_session_ended,,,,
U,2007/08/30 11:36:21:000 MST,CE,1,IBM.2107-75R0830,N,8020,
Web_SM_session_start,Web_SM_session_started,,,
```

```
U,2007/08/30 11:38:37:000 MST,CE,1,IBM.2107-75R0830,N,8022,
Web_SM_session_ended,,,,
U,2007/09/04 16:19:36:000 MST,,1,IBM.2107-75R0830,N,8022,
Web_SM_session_ended,,,,
U,2007/09/04 16:22:57:000 MST,hscroot,1,IBM.2107-75R0830,N,8020,
Web_SM_session_start,Web_SM_session_started,,,
U,2007/09/04 16:23:00:000 MST,hscroot,1,IBM.2107-75R0830,N,8022,
Web_SM_session_ended,,,,
```

Audit Log file definitions

Fields are output in comma-separated (CSV) format. This format makes it easy to import the file into a spreadsheet. The Input Parameters field is a special case. It uses the CSV format internally to separate one input field from the next. To manage this, the entire Input Parameters field is enclosed in double quotation marks.

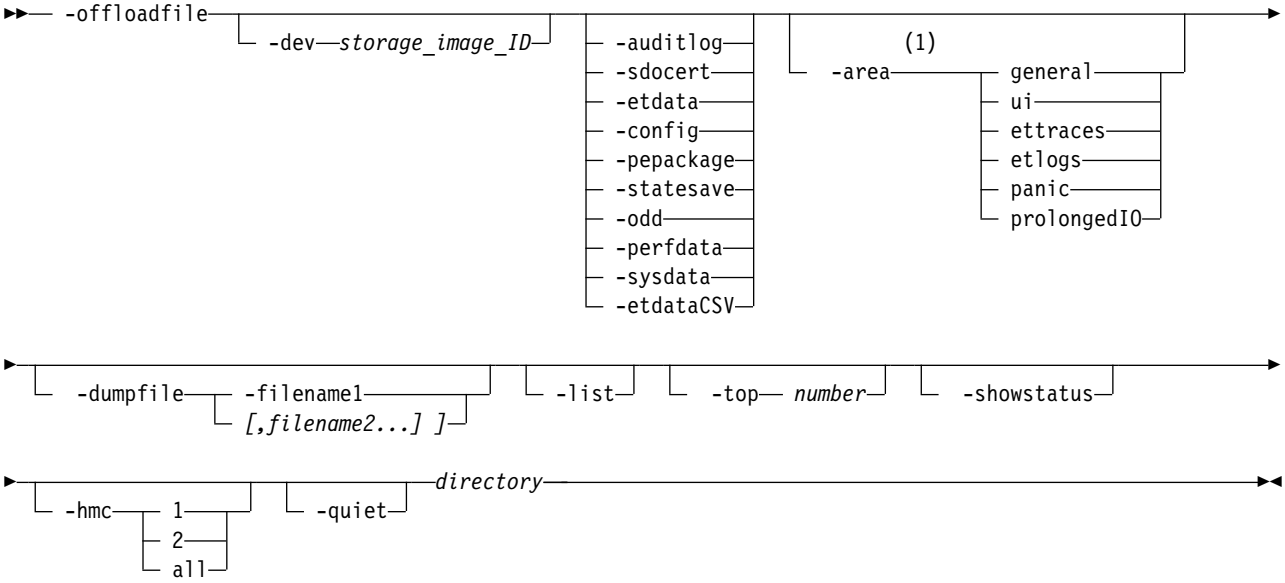
Field	Format	Description
Source	1 char	Specifies the source of the log entry: • **S** - Represents a server event that is not associated with a user action. • **U** - Represents a user-requested action. • **C**- Represents a continuation line for additional input attributes. There can be multiple **C** entries for a given user-requested (U) log entry.
Timestamp	YYYY/MM/DD HH:MM:SS:MMM TMZ	Represents the date, time, and time zone of the log entry.
User	1 - 16 char	Represents the user account that is making the request.
MC	1 char, a "1" or "2"	Represents the management console that processed the user request.
Device	16 char	Represents the storage image ID that consists of the following values: manufacture, type, and serial number.
NWC	1 char	Represents the following message types: **N** = notification, **W** = warning, and **C** = critical.
Entry ID	4 char	Represents the unique identifier that is associated with the activity that is represented by the log entry.
Entry name	20 char max	A text description that corresponds to the Entry ID.
Object ID	5 char max	Represents a unique identifier that identifies the object.
Exit code	8 char	Represents the final result code.

Field	Format	Description
Input Parameters	160 char max	Represents unformatted text that includes input parameters in the format: “attr1 = value1, attr2 = value2” with a comma (,) separator between parameters and double quotation marks around the entire field.

offloadfile

The **offloadfile** command exports the specified set of data files.

Note: The **offloadfile** command cannot be run from the Embedded DS CLI window.



Notes:

1 The **-area** parameter is only valid when the **-pepackage** parameter is specified.

Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, computer type, and serial number. The storage image ID is required if you do not specify a fully qualified storage image ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined dev id value for the current command.

-auditlog | -sdocert | -etdata | -config | -pepackage | -statesave | -odd | -perfdata | -sysdata | -etdataCSV

-auditlog
Exports one file that contains the audit log for the management console server.

-sdocert
Exports the SDO certificate for each storage facility.

-etdata
Exports two files that contain the IBM Easy Tier summary data.

-config
Offloads two files, one file that contains advanced settings, and another file that contains the Install Corrective Service settings.

Note: Only one file set parameter must be specified.

-pepackage
Generates a PE package and offloads it for use by product support teams, but only if you do not specify **-showstatus**. Specify only one file set.

-statesave
Exports a file set and offloads it for use by product support teams, but only if you do not specify **-showstatus** or **-list**. Specify only one file set.

-odd
Exports a file set and offloads it for use by product support teams, but only if you do not specify **-showstatus** or **-list**. Specify only one file set.

-perfdata
Downloads the performance summary CSV file.

-sysdata
Downloads the system summary file. Specify only one file set.

-etdataCSV
Downloads a zipped file that contains one or more Easy Tier CSV files and an Excel tool that is used for the CSV files. (The CSV files are Microsoft Excel files.) The zipped file is downloaded to a directory that you specify. Specify only one file set.

-area general | ui | ettraces | etlogs | panic | prolongedIO
(Optional) Specifies the following available options.

Note: The **-area** parameter is only valid when the **-pepackage** parameter is specified.

general
Offloads a general PE package to product support teams.

ui Sends an ESSNI/DSGUI/API PE package to product support teams.

ettraces
Sends Easy Tier debug traces to product support teams.

etlogs
Sends an automatic mode performance log to product support teams.

panic
Sends a panic log to product support teams.

prolongedIO
Sends long I/O traces to product support teams.

Note: A general package is generated and offloaded when you do not choose a parameter.

-dumpfile filename1 ***[,filename2 ...]***
(Optional) Specifies the dump files that are offloaded. When you do not specify the parameter, the most current file is offloaded.

-list
(Optional) Lists the dump files on both LPARs that you can offload. An error message displays when you use the **-list** parameter but do not specify the **-statesave** or **-odd** parameters.

-top ***number***
(Optional) Specifies the number of the dump files to display. An error message displays when you use the **-top** *number* parameter but do not specify the **-statesave** or **-odd** parameters.

-showstatus
(Optional) Use this parameter only with **-pepackage**, **-statesave**, or **-odd**. An error message displays when you use the **-showstatus** parameter but do not specify the **-statesave** or **-odd** parameters.

-hmc 1 | 2 | all
(Optional) Specifies the HMC on which you want to export files from. "-hmc 1" specifies the primary HMC, and "-hmc 2" specifies the secondary HMC. The default value "all" specifies the primary HMC on a single HMC system, and specifies both the primary and secondary HMCs on a dual HMC system.

If **-pepackage** is specified, "-hmc all" is not allowed. If the "-hmc" parameter is not specified, the default value is the primary HMC(HMC1). This option is ignored if "statesave" or "-odd" is specified.

-quiet
(Optional) Turns off the confirmation prompt for replacing an existing file.

directory
(Required) Specifies the local directory path that is used as the destination for offloading files.

Note: The User Access Control (UAC) settings for Windows Vista and later, or Windows Server 2008 and later, might not allow for the export files. (that use the **offloadfile** command) to a directory that requires elevated privileges. Unfortunately, the Windows operating system returns success and the **offloadfile** command show a message that states that the files were exported successfully, but the files do not exist in the specified directory. To work around this problem, you use one of the following options.

- Select a different directory that does not require elevated privileges to create a file.
- Click the DSCLI desktop shortcut and select **Run as Administrator**.

Examples

Example 1

Start the **offloadfile** command to export Easy Tier heat data.

```
dscli> offloadfile
-dev IBM.2107-75FA120 -etdata C:\temp
```

The resulting output

```
Sun Apr 09 02:23:49 PST 2004 IBM DS CLI
Offloadfile - The etdata file is offloaded to c:\temp\SF1300860ESS01_heat.data.
Offloadfile - The etdata file is offloaded to c:\temp\SF1300860ESS11_heat.data.
```

Example 2

Start the **offloadfile** command to export an SDO certificate.

```
dscli> offloadfile -dev IBM.2107-75FA120 -sdocert C:\temp
```

The resulting output

```
Date/Time - February 17, 2017 3:32:24 PM CST IBM DSCLI Version 0.0.0.0
DS:IBM.2107-75FA120
CMUC00428I offloadfile - File for sdocert is successfully
copied to C:\temp\SDOCertificate_2107-941-1300860_2011-01-26_09_37_54.
```

Example 3

Start the **offloadfile** command to export audit logs

```
dscli> offloadfile
-dev IBM.2107-75FA120 -auditlog C:\temp
```

The resulting output

```
Date/Time - February 17, 2017 3:32:24 PM CST DSCLI Version 0.0.0.0
DS - IBM.2107-75FA120
CMUC00428I offloadfile - File for auditlog is
successfully copied to c:\temp\IBM.-2107-75FA120_auditlog_HMC1.log.
```

Example 4

Start the **offloadfile** command to offload a general package for use by product support teams.

```
dscli> offloadfile
-dev IBM.2107-1300861 --pepackage
```

The resulting output

```
DS: IBM.2107-1300861

CMUC00497I offloadfile: Offloading request has been successfully submitted.
The system will complete the request momentarily.
```

Example 5

Start the **offloadfile** command to display the status of the PE package.

```
dscli> offloadfile
-dev IBM.2107-1300861 --pepackage -showstatus
```

The resulting output

```
DS: IBM.2107-1300861
CMUC00492I offloadfile: Offload task is running.
```

Example 6

Start the **offloadfile** command to offload config files.

```
dscli> offloadfile
-dev IBM.2107-1300861 -config -hmc 1 c:\temp
```

The resulting output

```
CMUC00428I offloadfile: The config file has been offloaded to C:\temp\productSettings.
CMUC00428I offloadfile: The config file has been offloaded to C:\temp\ICSInstall.history.
```

Example 7

Start the **offloadfile** command to offload the performance summary CSV file.

```
dscli> offloadfile
-dev IBM.2107-1300861 -perfdata c:\temp
```

The resulting output

```
CMUC00428I offloadfile: The perfdata file has been offloaded to
C:\temp\FORMATTED_PERF_SAMPLES_1300861.csv.
```

Example 8

Start the **offloadfile** command to offload the Easy Tier CSV files.

```
dscli> offloadfile
-dev IBM.2107-1300861 -etdataCSV C:/temp
```

The resulting output

```
CMUC00428I offloadfile: The etdataCSV file has been offloaded to C:/temp/ et_data_20170319083810.zip.
```

Example 9

Start the **offloadfile** command to offload the system summary file.

```
dscli> offloadfile
-dev IBM.2107-1300861 -sysdata C:/temp
```

The resulting output

```
CMUC00428I offloadfile: The sysdata file has been offloaded to
C:/temp/ 2107IBM.-1300861_StgSysSummary_20170317013659.csv.
```

Physical resource information commands

Commands are provided to view information about the physical resources in your storage complex.

The following physical resource information commands are available:

lsda Displays a report that lists the device adapters (DA) for each storage image and the status of each device adapter in the list.

lsddm Displays a report that lists the disk drive modules and status information for each disk drive module in the list.

lsframe
Generates a report that displays a list of the frame enclosures for a storage image.

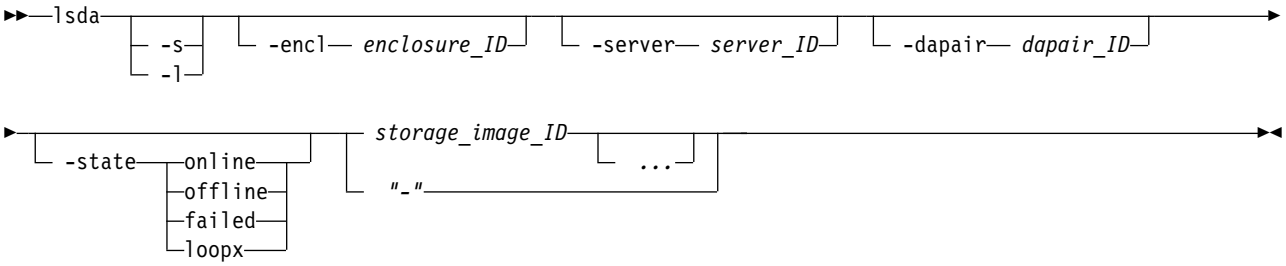
lshba Displays a report that lists the storage image host adapters and status information for each host adapter in the list.

lsserver
Displays all the servers in a storage complex or a list of specified servers. The displayed list also provides the status information for each server including the LIC version number, operating system version number, and bundle version.

lsstgencl
Generates a report that displays a list of the storage enclosures and status information for each enclosure in the list.

lsda

The **lsda** command displays a list of device adapters (DA) for each storage image. You can use this command to look at the status of each device adapter in the list.



Parameters

-s
(Optional) Displays only device adapter IDs. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output, plus the I/O enclosure and device adapter locations, feature codes, and interface IDs. You cannot use the **-l** and the **-s** parameters together.

-encl *enclosure_ID*
(Optional) Displays the device adapters that are associated with the specified processor node or I/O enclosure.

-server *server_ID*
(Optional) Displays only device adapters that are associated with the specified server.

-dapair *dapair_ID*
(Optional) Displays only device adapters that are associated with the specified device adapter pair.

-state online | offline | failed | loopx
(Optional) Displays only device drivers in the specified state.

storage_image_ID ... | -
(Required) Displays device adapters for the specified storage images. A storage image ID includes a value for the manufacturer, machine type, and serial number. You must separate multiple IDs with spaces. The ellipsis (...) indicates that, optionally, you can specify multiple values.

Note: You cannot specify ID ranges.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lsda** command.

Invoking the lsda command

```
dscli> lsda -l
IBM.2107-75FA120
```

The resulting output

ID	State	Loc
IBM.2107-75FA120/R1-11-P1-C1	Online	U2107-75FA120-11/P1-C1
IBM.2107-75FA120/R1-11-P1-C6	Online	U2107-75FA120-11/P1-C6
IBM.2107-75FA120/R1-12-P1-C1	Online	U2107-75FA120-12/P1-C1
IBM.2107-75FA120/R1-12-P1-C6	Online	U2107-75FA120/R1-12-P1-C6

				DA-Pair	
FC	Server	DA pair	Interfs	Type	WWN
1224	00	IBM.2107-75FA120/11	111,0112,0113,0114	2	500507604AC028E6
1234	01	IBM.2107-75FA120/12	111,0112,0113,0114	2	500507604AC028E7
1234	00	IBM.2107-75FA120/11	111,0112,0113,0114	2	500507604AC028E8
1234	01	IBM.2107-75FA120/12	111,0112,0113,0114	2	500507604AC028E9

Report field definitions

ID*

Indicates the unique identifier of the device adapter.

State

Indicates the current availability state of the specified device adapter. One of the following values is displayed:

- Online
- Offline
- Coming online
- Going offline
- Failed
- Offline Loop 1
- Offline Loop 2
- Offline Loop 3
- Offline Loop 4
- Offline Loop 1/2
- Offline Loop 3/4
- Taking down Loop 1
- Taking down Loop 2
- Taking down Loop 3
- Taking down Loop 4
- Taking down Loop 1/2
- Taking down Loop 3/4
- Bring up all loops

Loc+

Indicates the I/O enclosure and the device adapter location.

The I/O enclosure location format is *Uttt.mmm.ppssssss*.

The device adapter location format is *Pn-Cn* where *Pn* indicates the Planner number (1) and *Cn* indicates the card number (1 - 6).

FC+

Indicates the feature code that is used to order the specified device adapter.

Server

Indicates the server or device adapter group to which the device adapter is assigned.

DA pair

Indicates the storage unit ID that is followed by the device adapter pair ID that is associated with the specified device adapter. The device adapter pair identifier is a two-digit decimal number, with no leading zeros. Device adapter pairs are located in I/O enclosure pairs. Device adapter pair ID implies I/O enclosure location.

An even-numbered device adapter pair ID indicates the first device adapter pair in an I/O enclosure pair. An odd-numbered device adapter pair ID indicates the second device adapter pair in an IO enclosure pair.

Interfs+

Indicates the four interface IDs that are associated with the FC-AL ports.

DA_Pair_Type+

Indicates the DA Pair Type and the enclosure type.

0 Unknown

1 Indicates an enclosure that holds up to 16 LFF drives.

2 Indicates an enclosure that holds up to 24 SFF or 12 LFF drives.

3 Indicates an enclosure that holds up to 30 high-performance flash devices.

4 Indicates that one pair of high-performance flash enclosure Gen2 enclosures can hold up to 48 high-performance flash devices.

WWN+

Indicates the worldwide Name of the DA/enclosure (example 500507604AC028E6).

Displays a '-' (dash) if the value is unknown.

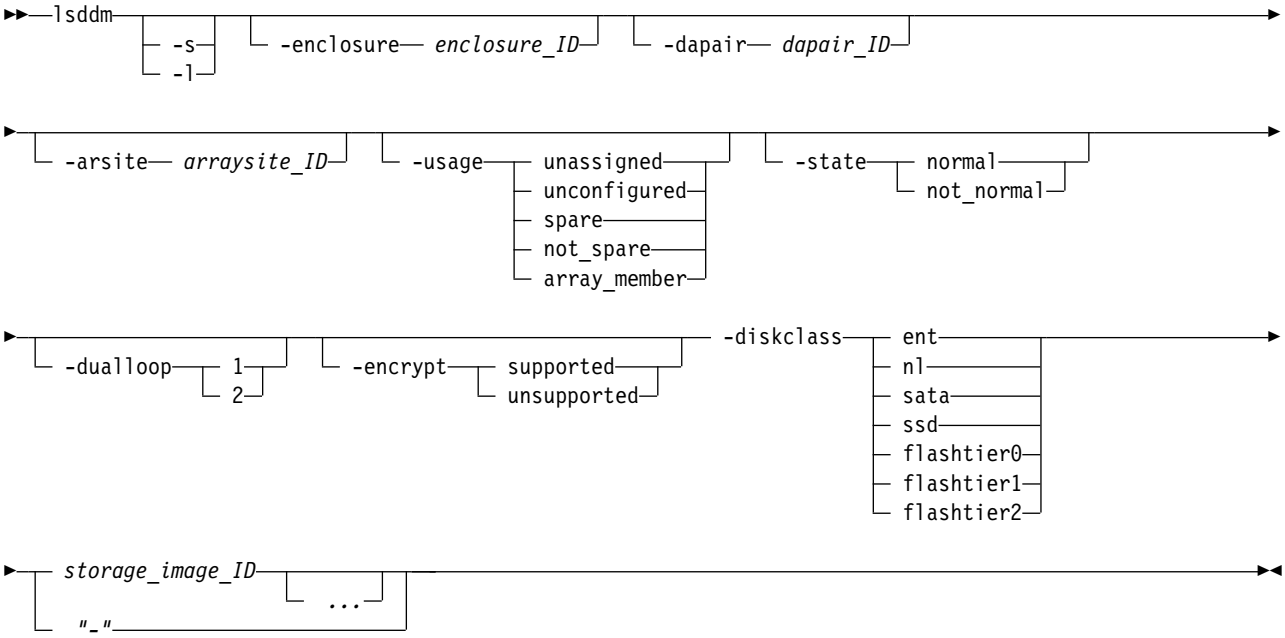
Key:

***** Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

lsddm

The **lsddm** command displays a list of disk drive modules (DDMs) and status information for each DDM in the list.



Parameters

-s (Optional) Displays only the DDM IDs. You cannot use the **-s** and **-l** parameters together.

-l (Optional) Displays the default output. You cannot use the **-s** and **-l** parameters together.

-enclosure *enclosure_ID*

(Optional) Specifies that the system displays DDMs that are associated with a common storage enclosure ID. This parameter accepts a fully qualified enclosure ID, which includes either the storage image ID or a shortened version without the storage image ID. The shortened version is a hexadecimal number within the range (00 - 3F).

-dapair *dapair_ID*
(Optional) Specifies that the system displays DDMs that are associated with a common device adapter (DA) pair ID. This parameter accepts a fully qualified DA pair ID, which includes either the storage image ID or a shortened version without the storage image ID. The shortened version is a two-digit decimal number with no leading zeros.

-arsite *arraysite_ID*
(Optional) Specifies that the system displays DDMs that are associated with a common array site ID. This parameter accepts a fully qualified array site ID, which includes either the storage image ID or a shortened version without the storage image ID. The shortened version is a four-digit decimal number with no leading zeros, prefixed with the letter *S*.

-usage unassigned | unconfigured | spare | not_spare | array_member
(Optional) Specifies that the system displays DDMs that are associated with a specified usage.

-state normal | not_normal
(Optional) Specifies that the system (with one exception) displays DDMs that are associated with a specified state. Failed DDMs that are deferred service and are associated with state=10 no longer display.

-dualloop 1 | 2
(Optional) Specifies that the system displays DDMs that are associated with the designated dual loop.

-encrypt supported | unsupported
(Optional) Specifies that the system displays only the DDMs that have the specified encryption capability.

-diskclass ent | nl | sata | ssd | flashtier0 | flashtier1 | flashtier2
(Optional) Displays the DDMs that are associated with the specified disk class.

sod Specifies solid-state devices.

nl Specifies high-capacity near-line disk drives.

ent Specifies high-speed Enterprise disk drives.

data Specifies high-capacity SATA devices.

flashtier0
Specifies high performance flash (Flash Tier 0) drives.

flashtier1
Specifies the first tier of high capacity flash (Flash Tier 1) drives.

flashtier2
Specifies the second tier of high capacity flash (Flash Tier 2) drives.

storage_image_ID **... | -**
(Required) Specifies that the system displays DDM information for the designated storage image IDs. A storage image ID includes a value for the manufacturer, machine type, and serial number. You can specify multiple IDs and they must be separated with a space between each ID. The ellipsis (...) indicates that, optionally, you can specify multiple values.

Note: You cannot specify ID ranges.

If you use the dash (-), the specified value is read from standard input.

Note: You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 1

For this command, the results are shown in table format to provide clarity. The actual reports do not display as tables. See Example 3 for output descriptions.

```
dscli> lsddm -enclosure IBM.2107-D02-011GP
```

ID	DA pair	Dkcap (10^9B)	dkuse	arsite	State
IBM.2107-D02-011GP/ R1-P1-D1	0	146.0	array member	S6	Normal
IBM.2107-D02-011GP/ R1-P1-D2	0	146.0	array member	S6	Normal
IBM.2107-D02-011GP/ R1-P1-D3	0	146.0	array member	S3	Normal
IBM.2107-D02-011GP/ R1-P1-D4	0	146.0	array member	S4	Normal
IBM.2107-D02-011GP/ R1-P1-D5	0	146.0	array member	S5	Normal

Example 2

For this command, the results are shown in table format to provide clarity. The actual reports do not display as tables. See Example 3 for output descriptions.

```
dscli> lsddm
```

ID	DA Pair	dkcap (10^9B)	dkuse	array	State
IBM.2107-D04-460AY/ R1-Q1-E1	9	7688.0	array member	S3	Normal
IBM.2107-D04-460AY/ R1-Q1-E1	9	7688.0	array member	S5	Normal

Example 3

```
dscli> lsddm -l IBM.2107-75FA120
```

ID	Model	Loc	Firmware level	DA pair	DualLoop	Dkcap (10^9B)
IBM.2107-D21-75FA120/ R1-S11-P1-D1	S0BE146	U2107.921. 75FA120	3603	11	1	145
IBM.2107-D21-75FA120/ R1-S11-P1-D2	S0BE146	U2107.921. 75FA120	3603	11	1	145

Diskrpm	Dkinf	Dkrate (Gb/sec)	Dkuse	Arsite	Position	State	Diskclass	DDMSN	WWNN	Encrypt	Encryption
1000	SAS	2	array member	S123	1	normal	ENT	0123456789 ABCDEF	3037AAF30 X23990X	supported	encrypted
1000	SAS	2	array member	S123	2	normal	SATA	0123456789 ABCDEF	3037AAF30 X23990X	supported	encrypted

Report field definitions

ID*

Indicates the system-assigned unique identifier for the DDM.

Model

Indicates the DDM model. The model name is a string of the form *VRFCGGG*, where *VRFC* is the type of disk family and *GGG* is the disk capacity in decimal gigabytes (GB).

Loc

Indicates the storage enclosure and the DDM location. The DDM location format is *Pn-Dn*, where *Pn* is the Planer number (1), and *Dn* is the DDM number (1 - 16).

Firmware level

Indicates the level of firmware that is installed on the specified DDM.

DA pair
Indicates the device adapter pair ID. DA pairs are in I/O enclosure pairs.

Note: An even-numbered DA pair ID indicates the first DA pair in an I/O enclosure pair. An odd-numbered DA pair ID indicates the second DA pair in an I/O enclosure pair.

Dualloop
Indicates the dual loop where the specified DDM is stored. The value is either 1 or 2.

Dkcap (10^9B)
Indicates the DDM raw capacity in decimal gigabytes (GB).

Diskrpm
Indicates the DDM rpm. One of the following values is displayed:

- 10000
- 15000

Dkinf
Indicates the DDM interface type:

- SAS (Serial Attached SCSI)

Dkrate (Gb/sec)
Indicates the DDM interface rate.

Dkuse
Indicates the DDM usage in an array site. One of the following values are displayed:

Unassigned
The storage device is not assigned to an array site. The array site and array site position attributes are not valid.

Unconfigured
The storage device is assigned to an array site, but the array site is not configured and the storage device is not a required spare.

Spare required
The storage device is assigned to an array site and is a required spare.

Spare not required
The storage device is assigned to an array site that is configured with an array and a spare. However, the spare is not a required spare. You might reconfigure this array site from a narrow array to a wide array.

Array member
The storage device is assigned to an array site that is configured and the array site position of the storage device makes it part of the array.

Arsite
Indicates the array site ID.

Position
Indicates the DDM position in an array site configuration of DDMs.

State
Indicates the current DDM state. One of the following values is displayed:

Normal
The storage device is operational and functional in its current disk usage.

New
Indicates the initial state when a DDM is inserted or first discovered.

Installing
A new storage device that was identified.

Verifying
The storage device is made accessible to the device adapter. The characteristics are determined, cabling is checked, and diagnostics are run.

Formatting
A verified storage device requires low-level formatting and the formatting operation is in progress.

Initializing
The storage device is being initialized with all zero sectors.

Certifying
The storage device is read-accessed to determine that all sectors can be read.

Rebuilding
The storage device is being rebuilt with data from the array that it is associated with.

Migration Target
DDM migration is migrating another array member storage device to this spare storage device.

Migration Source
DDM migration is migrating this array member storage device to another spare storage device.

Failed
The storage device failed and an immediate repair action is required.

Failed - Deferred Service
The storage device failed and a repair action is not immediately required.

Removed
The storage device is removed from the system and removal was processed by the system.

Inappropriate
The storage device is incompatible with the system; for example, a storage device that has the wrong capacity or rpm. The DDM is not failed because it can be valid for other systems and locations.

Inter failed
Indicates that the DDM is faulty but still working.

PFSed
Indicates that the DDM is prepared for service, and ready to be removed without impacting the system.

Diskclass
Indicates the disk class. One of the following values can be displayed:

ENT
Indicates enterprise and represents high-speed Fibre Channel disk drives.

NL Indicates near-line and represents ATA (FATA) disk drives.

SATA
Indicates high-capacity SATA disk drives.

SSD
Indicates solid-state devices.

FlashTier0
Indicates high performance flash (Flash Tier 0) drives.

FlashTier1
Indicates the first tier of high capacity flash (Flash Tier 1) drives.

FlashTier2
Indicates the second tier of high capacity flash (Flash Tier 2) drives.

DDMSN+
Indicates the list of DDM serial numbers (SN) that are associated with the designated array site. Each DDM SN is a 16-character string. Each serial number is separated by a comma.

WWNN+
Indicates the worldwide node name that is assigned to the storage image.

Encrypt+
Indicates the encryption support capability. One of the following values can be displayed:

Supported
The disk drive modules in this rank support encryption.

Unsupported
The disk drive modules in this rank do not support encryption.

Encryption
Indicates whether specified drives are encrypted.

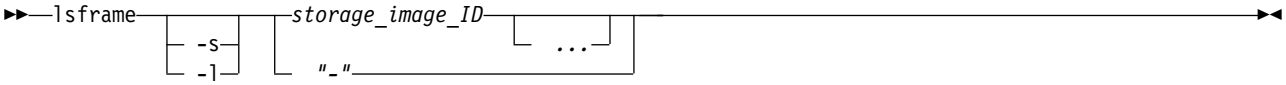
Key:

* Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

lsframe

The **lsframe** command displays a list of frame enclosures for a storage image.



Parameters

-s Displays the rack enclosure ID. You cannot use the **-l** and the **-s** parameters together.

-l Displays default output plus the frame ID and location of the frame enclosure. You cannot use the **-l** and the **-s** parameters together.

storage_image_ID ... | **-**
Displays frame enclosure information for the specified storage image IDs. A storage image ID includes manufacturer, type, and serial number. You must separate multiple IDs with spaces. The ellipsis (...) indicates that, optionally, you can specify multiple values.

Note: ID ranges cannot be specified.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DSCLI interactive command mode.

Example

For this command and all other DSCLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lsframe** command.

Invoking the lsframe command

```
dscli> lsframe -l IBM.2107-75FA120
```

The resulting output

ID	frm#	Frame	loc
IBM.2107-75FA120/R1	R1	1	U2107-75FA120
IBM.2107-75FA120/R2	R2	2	U2107-75FA120

Report field definitions

ID*
Identifies the unique identifier of the frame enclosure.

frm#
Identifies the frame number within a storage system that contains the specified frame enclosure.

Frame+
Identifies the unique identifier of the storage system equipment frame that contains the specified frame enclosure.

Loc+
Identifies the frame enclosure location.

The location format is *Uttt.mmm.ppssssss*.

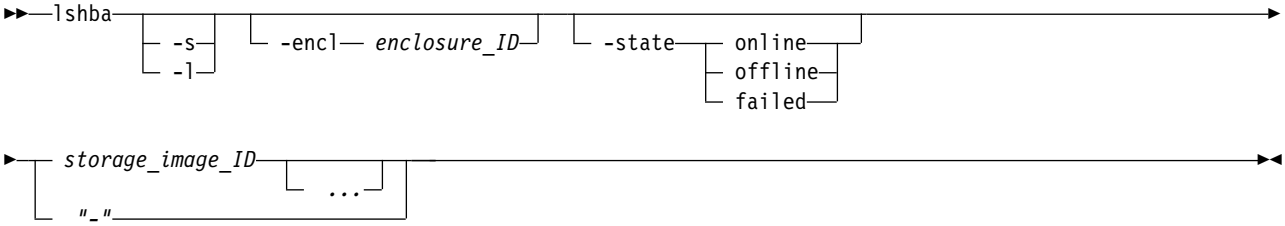
Key:

* Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

lshba

The **lshba** command displays a list of storage image host adapters and status information for each host adapter in the list.



Parameters

-s (Optional) Displays only the host adapter IDs. You cannot use the **-l** and the **-s** parameters together.

-l (Optional) Displays the default output plus the host adapter feature code and interface IDs. You cannot use the **-l** and the **-s** parameters together.

-encl *enclosure_ID*
(Optional) Specifies that the system displays host adapters that are associated with a common processor node or I/O enclosure ID.

-state online | offline | failed
(Optional) Specifies that the system displays host adapters that are in a specified state.

storage_image_ID ... | -
(Optional) Specifies that the system displays host adapter information for the designated storage image IDs. A storage image ID includes a value for the manufacturer, machine type, and serial number.

Notes:

1. Multiple IDs must be separated with spaces.
2. ID ranges cannot be specified.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input.

Note: You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

The following example represents the headers that are displayed on the output report that is associated with the **lshba** command.

Invoking the lshba command

```
dscli> lshba -l IBM.2107-75FA120
```

The resulting output

```
ID                               State            Loc
=============================================================
IBM.2107-75FA120/R1-11-P1-C2     Online   U2107-75FA120-11/P1-C2
IBM.2107-75FA120/R1-11-P1-C3     Online   U2107-75FA120-11/P1-C3
IBM.2107-75FA120/R1-12-P1-C2     Online   U2107-75FA120-/R1-12-P1-C2
IBM.2107-75FA120/R1-12-P1-C3     Online   U2107-75FA120/R1-12-P1-C3

FC   interfID
================================
1234 0x0121,0x0121,0x0123,0x0124
1234 0x0131,0x0131,0x0133,0x0134
1234 0x0221,0x0221,0x0223,0x0224
1234 0x0231,0x0231,0x0233,0x0234
```

Report field definitions

ID Indicates the unique identifier of the host adapter.

State
Indicates the current availability state of the specified host adapter.

Loc
Indicates the I/O enclosure and the host adapter location.

The I/O enclosure location format is *Uttt.mmm.ppssss*.

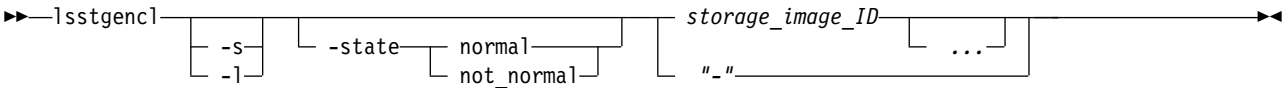
The host adapter location format is *Pn-Cn* where *Pn* indicates the Planner number (1) and *Cn* indicates the card number (1 - 6).

FC Indicates the feature code that is used to order the specified host adapter.

interfID
Indicates the four interface IDs that are associated with the I/O ports on the host adapter.

lsstgencl

The **lsstgencl** command displays a list of storage enclosures and status information for each enclosure in the list.



Parameters

-s (Optional) Displays the storage enclosure ID. You cannot use the **-l** and the **-s** parameters together.

-l (Optional) Displays default output, plus the frame number, ID, location, feature code, number of storage slots, and state of the storage enclosure. You cannot use the **-l** and the **-s** parameters together.

-state normal| not_normal
(Optional) Displays all the storage enclosures that are associated with the specified storage system that contain a condition of normal or a condition that falls under the category of not normal.

storage_image_ID ... | -
(Required) Displays storage enclosure information for the specified storage image IDs. A storage image ID consists of manufacturer, machine type, and serial number. You must separate multiple IDs with a space between each ID. The ellipsis (...) indicates that, optionally, you can specify multiple values.

Note: You cannot specify an ID range.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lsstgencl** command by using the **-l** parameter.

Invoking the lsstgencl command

```
dscli> lsstgencl -l IBM.2107-75FA120
```

The resulting output

ID	Frame	Enclnum	Loc	FC	Interfaces
IBM.2107-D21.75FA120/R1-S11	R1	S11	U2107.921.75FA120	3221	0111, 0112, 0113, 0114
IBM.2107-D21.75FA120/R1-S12	R1	S12	U2107.921.75FA120	3221	0121, 0122, 0123, 0124

Interadd	Storslot	Stordev	Cap (GB)	RPM	State	DA_Pair_Type
0 x 0	16	16	145	10000	normal	1
0 x 1	16	16	145	10000	normal	1

Report field definitions

ID*
Indicates the enclosure ID and enclosure number.

Note: This is the only information that is displayed if you use the **-s** parameter. None of the other values are displayed.

Frame+
Indicates the frame number within the designated storage system that contains this storage enclosure.

Enclnum+
Indicates the identity of a storage enclosure within a storage system frame.

Loc+
Indicates the storage enclosure location by identifying the storage system frame that contains the storage enclosure. The location format is *Utttt.mmm.ppssssss*.

FC+
Indicates the feature code that was used to order this storage enclosure.

Interfaces
Indicates a list of four interface identifiers, one for each DA I/O port.

A DA interface ID consists of four hexadecimal characters with the following format: *t00 eeeee aaaa pppp*, (value is separated for readability), where

- *t* = port type (0 = I/O port, DA ports are always I/O ports)
- *00* = reserved
- *eeeee* = enclosure number
- *aaaa* = adapter number
- *pppp* = port number

Notes:

1. For dual loop 0 enclosures, the DA I/O port values are displayed as 0080, 0081, 0180, 0181.
2. For dual loop 1 enclosures the DA I/O port values are displayed as 0082, 0083, 0182, 0183.

Interadd
Indicates the FC-AL interface base address that is assigned to this storage enclosure for DDM access.

Storslot+
Indicates the number of slots for storage devices in this storage enclosure.

Stordev
Indicates the number of storage devices that are installed in this storage enclosure.

Cap (GB)
Indicates the capacity of DDMs in the storage enclosure.

Note: This field can contain multiple capacity values that are separated by a comma when the DDMs with different capacity are installed in the storage enclosure.

RPM
Indicates the rpm of the DDMs in the storage enclosure.

Note: This field can contain multiple RPM values that are separated by a comma when the DDMs with different capacity are installed in the storage enclosure.

State+
Indicates the condition of the storage enclosure. The condition of the enclosure is either **normal** or one of the conditions that falls under the category **not_normal**. The following values can be displayed:

normal
Indicates that the storage enclosure is not in any failure or transient condition.

offline
Indicates that the storage enclosure is not capable of processing any functions.

failed
Indicates that the storage enclosure is broken and ready to be removed without impacting the system.

resuming
Indicates that the storage enclosure is in the process of coming online.

quiescing
Indicates that the storage enclosure is in the process of going offline.

DA_Pair_Type+
Indicates the DA Pair Type and the enclosure type, as follows:

0 Unknown

1 Indicates an enclosure holding up to 16 LFF drives.

2 Indicates an enclosure holding up to 24 SFF or 12 LFF drives.

3 Indicates an enclosure holding up to 30 high-performance flash devices.

4 Indicates that one pair of high-performance flash enclosure Gen2 enclosures can hold up to 48 high-performance flash devices.

Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

Remote event notification commands

Use the remote event notification commands to forward logs automatically and securely to a central repository for auditing.

The following remote event notification commands are available:

chsyslogserver
Modifies a Syslog server.

lssyslogserver
Displays a list of Syslog servers.

mksyslogserver
Creates a Syslog server.

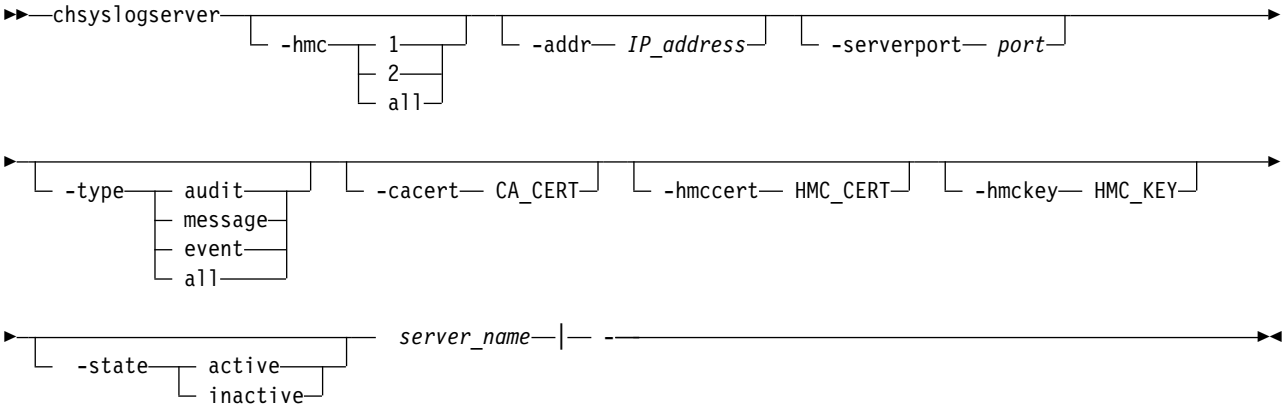
rmsyslogserver
Removes a Syslog server.

chsyslogserver

The **chsyslogserver** command modifies a Syslog server.

Syslog servers for each HMC can be created separately. For the parameter **-hmc**, the default value all specifies the primary HMC on a single HMC system, and specifies both the primary and secondary HMCs on a dual HMC system. If any of HMCs fails, an error is returned and rollback is done on both HMCs.

Note: The **chsyslogserver** command cannot be run from the Embedded DS CLI window.



Parameters

-hmc 1 | 2 | all
(Optional) Specifies the HMC on which you want to modify the server.

- **-hmc 1** specifies the primary HMC
- **-hmc 2** specifies the secondary HMC.
- The default value **all** specifies the primary HMC on a single HMC system, and specifies both the primary and secondary HMCs on a dual HMC system.

-addr *IP_address*
(Optional) Specifies the TCP/IP address for the Syslog server.

-serverport *port*
(Optional) Specifies the communications port of the Syslog server. If you do not specify this parameter, the port of the Syslog server is not changed.

-type audit| message | event | all
(Optional) Specifies an array of one or more types of logs to be forwarded to the Syslog server. You can specify one or more types of logs, which are separated by commas.

audit
Specifies that audit log information is forwarded to the Syslog server.

message
Specifies that system messages are forwarded to the Syslog server.

event
Specifies that system event information is forwarded to the Syslog server.

all
Specifies that all types of logs (audit, message, and event) are forwarded to the Syslog server. The default value is **all.**

-cacert CA_CERT
(Optional) Specifies the location for CA certificates. This option is only valid when the protocol is TLS.

-hmccert HMC_CERT
(Optional) Specifies the location for HMC certificates. This option is only valid when the protocol is Transport Layer Security (TLS).

-hmckey HMC_KEY
(Optional) Specifies the location for HMC private keys. This option is only valid when the protocol is TLS.

-state active|inactive
(Optional) Specifies whether the current server is active or not.

server_name | **-**
(Required) Specifies the Syslog server. Alternatively, the command accepts input from stdin when the dash (-) is specified.

Example

Invoking the chsyslogserver command

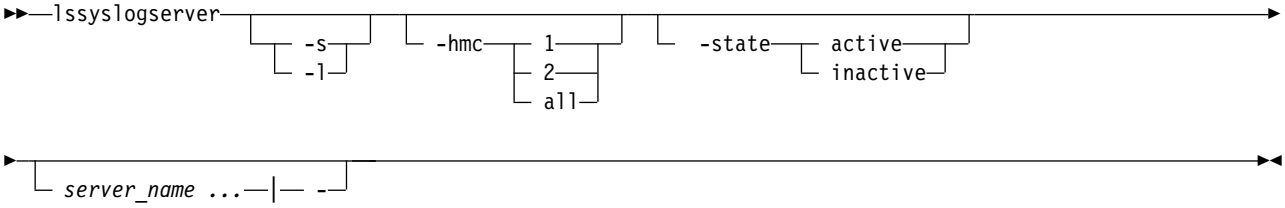
```
chsyslogserver -state active syslog_1
```

The resulting output

```
Syslog server syslog_1 successfully modified.
```

lssyslogserver

The **lssyslogserver** command displays a list of Syslog servers.



Parameters

-s | -l
(Optional) Specifies that a shorter version of the command output displays. The output includes the server name, IP address, state, and HMC. This value is the default.

Note: The **-s** and **-l** parameters cannot be used together.

The optional **-l** (long) parameter displays default output plus more attributes.

-hmc 1 | 2 | all
(Optional) Specifies the HMC that you want to query.

- **-hmc 1** specifies the primary HMC
- **-hmc 2** specifies the secondary HMC.
- The default value **all** specifies the primary HMC on a single HMC system, and specifies both the primary and secondary HMCs on a dual HMC system.

-state active|inactive
(Optional) Specifies whether the current server is active or not.

server_name| **-**
(Optional) Specifies the Syslog server. Alternatively, the command accepts input from stdin when the dash (-) is specified.

Example

Invoking the lssyslogserver command

```
dscli> lssyslogserver
```

The resulting output displays below

```
Fri Aug 17 17:23:49 PST 2017 IBM DSCLI Version 5.0.0.0 DS: IBM.2107-75A128
name      IP address      port    state      access    HMC

syslog1   9.100.75.7      514      active     online     2
syslog2   9.123.236.180   514     inactive    online     1
```

Example

```
dscli> lssyslogserver -l
```

The resulting output

```
Fri Aug 17 17:23:49 PST 2016 IBM DSCLI Version 5.0.0.0 DS: IBM.2107-75A128
name      IP address      port   state       access     type           protocol   HMC

syslog1   9.100.75.7      514    active      online     audit          tls        2
syslog2   9.123.236.180   514    inactive    online     event message  tls        1
```

Report field definitions

name
Indicates the name of the Syslog server.

IP_address
Indicates the IP address of the Syslog server.

port
Indicates the communication port of the Syslog server.

state
Indicates whether the Syslog server is active or not.

access
Indicates whether the Syslog server is accessible or not. Unknown displays when the information is not available.

protocol tcp | tls
Indicates the protocol type.

HMC
Indicates which HMC forwards its audit log to the Syslog server. The primary HMC is indicated by "1" and the secondary is indicated by "2". The output displays this value only when you specify the -l parameter.

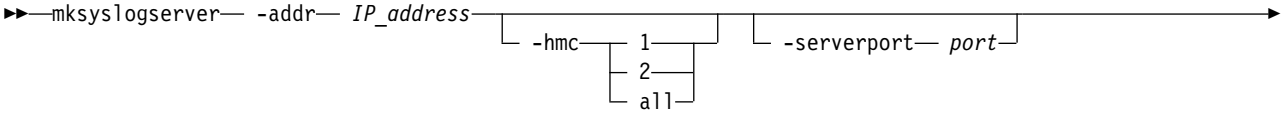
mksyslogserver

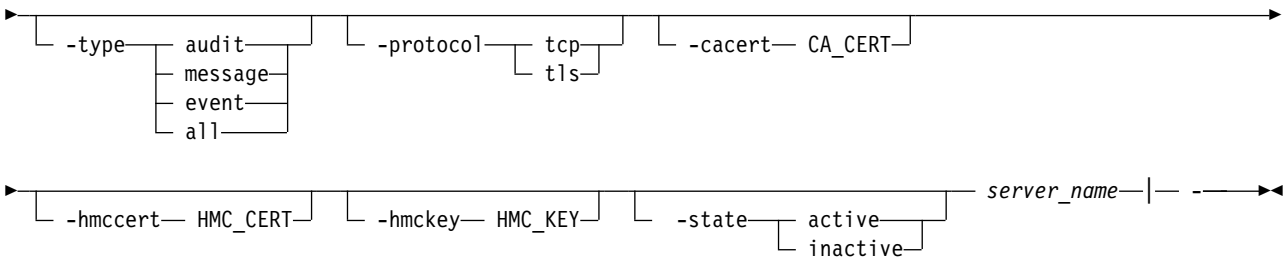
The **mksyslogserver** command creates a Syslog server.

Syslog servers for each HMC can be created separately. For the parameter **-hmc**, the default value all specifies the primary HMC on a single HMC system, and specifies both the primary and secondary HMCs on a dual HMC system. If any of HMCs fails, an error is returned and rollback is done on both HMCs.

Note: The **mksyslogserver** command cannot be run from the Embedded DS CLI window.

You can now forward one or more types of logs, including events and messages to the Syslog server.





Parameters

-addr *IP_address*
(Required) Specifies the TCP/IP address for the Syslog server.

-hmc 1 | 2 | all
(Optional) Specifies the HMC on which you want to forward audit logs.

- **-hmc 1** specifies the primary HMC
- **-hmc 2** specifies the secondary HMC.
- The default value **all** specifies the primary HMC on a single HMC system, and specifies both the primary and secondary HMCs on a dual HMC system.

-serverport *port*
(Optional) Specifies the communications port of the Syslog server. The default value is 514.

-type audit| message | event | all
(Optional) Specifies an array of one or more types of logs to be forwarded to the Syslog server. You can specify one or more types of logs, which are separated by commas.

audit
Specifies that audit log information is forwarded to the Syslog server.

message
Specifies that system messages are forwarded to the Syslog server.

event
Specifies that system event information is forwarded to the Syslog server.

all
Specifies that all types of logs (audit, message, and event) are forwarded to the Syslog server. The default value is **all**

-protocol tcp | tls
(Optional) Specifies the type of protocol. All Syslog servers that are configured for the storage system must use the same protocol. The protocols include the Transmission Control Protocol (TCP) specified with the **-tcp** parameter or the Transport Layer Security (TLS) specified with the **-tls** parameter.

If you specify TLS, all Syslog server configuration use the same following certificate files in all Syslog server configurations to ensure that Syslog messages are encrypted when processed on the server.

CA certificate file
Specified with the **-cacert CA_CERT** parameter.

HMC certificate file
Specified with the **-hmccert HMC_CERT** parameter.

HMC private key
Specified with the **-hmckey HMC_KEY** parameter.

Note: If this is the first Syslog server for the storage system and the protocol is TLS, the **-cacert**, **-hmccert**, and **-hmckey** parameters are required. If this Syslog server is not the first server defined for

the storage system and its protocol is TLS, then do not specify the **-cacert**, **-hmccert**, and **-hmckey** parameters. Otherwise, the parameters will replace the existing ones on the HMCs.

-cacert CA_CERT
(Optional) Specifies the location for the CA certificates. This option is only valid when the protocol is TLS.

-hmccert HMC_CERT
(Optional) Specifies the location for the HMC certificates. This option is only valid when the protocol is TLS.

-hmckey HMC_KEY
(Optional) Specifies the location for the HMC private keys. This option is only valid when the protocol is TLS.

-state active|inactive
(Optional) Specifies whether the current server is active or not.

server_name| **-**
(Required) Specifies the Syslog server. Alternatively, the command accepts input from stdin when the dash (-) is specified.

Example

Invoking the mksyslogserver command

```
dscli> mksyslogserver -ipaddr 9.110.75.7 -serverport 514 -type all syslog_1
```

The resulting output

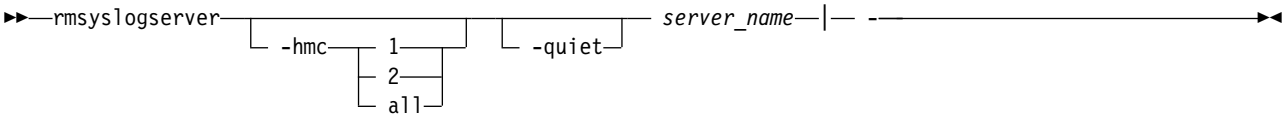
```
Date/Time: Sun Aug 11 02:23:49 PST 2016 IBM DSCLI Version: 5.0.0.0
DS: IBM.2107-75FA120
Syslog server syslog_1 successfully created.
```

rmsyslogserver

The **rmsyslogserver** command removes a Syslog server.

Note: A warning is displayed when the server to be deleted is active.

For parameter **-hmc**, the default value **all** specifies the primary HMC on a single HMC system, and specifies both the primary and secondary HMCs on a dual HMC system. If any of HMCs fails, an error is returned and rollback is done on both HMCs.



Parameters

-hmc 1 | 2 | all
(Optional) Specifies the HMC on which you want to delete a server.

- **-hmc 1** specifies the primary HMC
- **-hmc 2** specifies the secondary HMC.
- The default value **all** specifies the primary HMC on a single HMC system, and specifies both the primary and secondary HMCs on a dual HMC system.

-quiet
(Optional) Turns off the confirmation prompt for this command.

server_name| **-**
(Required) Specifies the Syslog server. Alternatively, the command accepts input from stdin when the dash (-) is specified.

Example

Invoking the rmsyslogserver command

```
dscli> rmsyslogserver syslog_1
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DSCLI Version: 5.0.0.0
DS: IBM.2107-75FA950 Syslog server syslog_1 successfully removed.
```

Resource Group commands

Commands that are used to specify policy-based limitations on the access and use of resources are referenced.

The following Resource Group commands are available:

chresgrp
Changes a resource group object on a storage image.

lsresgrp
Displays a list of resource group objects on the storage image.

manageresgrp
Manages the contents of a resource group object on a storage image.

mkresgrp
Creates a resource group object on a storage image.

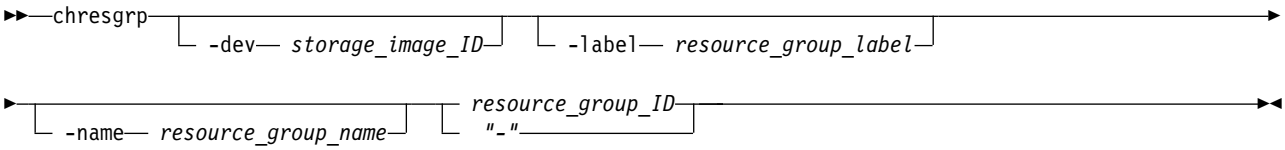
rmresgrp
Removes a resource group object on a storage image.

showresgrp
Displays detailed properties for an individual volume.

chresgrp

The **chresgrp** command is used to change a resource group object on a storage image.

Note: Resource Group 0 (zero) is predefined and cannot be created, deleted, or modified. By default, all resources belong to this group unless otherwise specified.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of the manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified resource group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-label *resource_group_label*
(Optional) Specifies the resource group label. The resource group label is 1 to 32 characters and is limited to upper and lower case alphabetic and numeric characters, and the special characters (-), (_), and (.). Label names must be unique.

-name *resource_group_name*
(Optional) Specifies the user-assigned nickname for this resource group object. The maximum length is 64 single-byte or 32 double-byte characters.

resource_group_ID | **-**
(Required) The resource group ID. The resource group ID begins with the letters *RG* and ends with a decimal number. If you use the dash (-), the specified value is read from standard input.

Example

Invoking the chresgrp command

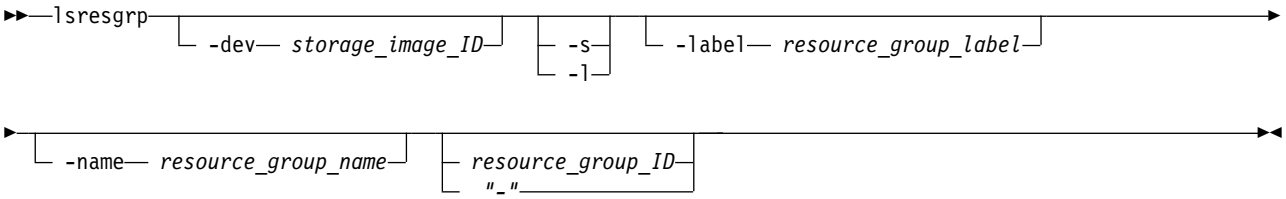
```
dscli> chresgrp -dev IBM.2107-75FA120 -name "A_Group" RG1
```

The resulting output

```
Resource Group RG1 successfully modified.
```

lsresgrp

The **lsresgrp** command displays a list of resource group objects on the storage image.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-s
(Optional) Displays only the resource group IDs. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays default output. You cannot use the **-l** and the **-s** parameters together.

-label *resource_group_label*
(Required) Displays the specified resource group label. The resource group label is 1 to 32 characters and is limited to upper and lower case alphabetic and numeric characters, and the special characters (-), (_), and (.).

-name *resource_group_name*
(Optional) Displays the specified user assigned nickname for this resource group object. The maximum length is 64 single-byte or 32 double-byte characters.

resource_group_ID | **-**
(Optional) The resource group ID. The resource group ID begins with the letters *RG* and ends with a decimal number. If the resource group ID is not specified, one will be assigned. If you use the dash (-), the specified value is read from standard input.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lsresgrp** command using the **-l** parameter.

Note:

Invoking the lsresgrp command

```
dscli> lsresgrp
```

The resulting output

ID	Name	State	Label
RG1	ProductA_admins	normal	Product_A
RG2	ProductA_group1	normal	Product_A.grp1
RG3	ProductA_group2	normal	Product_A.grp2

Report field definitions

ID*
Indicates the unique identifier that is assigned to this resource group ID.

name
Indicates the nickname that you assigned for this resource group object.

state
Indicates the current configuration state of this resource group. One of the following values is displayed:

Normal
Indicates that the resource group is not being configured.

Configuring
Indicates that the resource group is being configured.

Configuration Error
Indicates that the resource group configuration process failed to complete successfully.

Deconfiguring
Indicates that a resource group is in the process of being deleted.

Deconfiguration Error
Indicates that the resource group deletion process failed to complete successfully.

Label
Indicates the resource group label. The resource group label is 1 to 32 characters and is limited to upper and lower case alphabetic and numeric characters, and the special characters (-), (_), and (.).

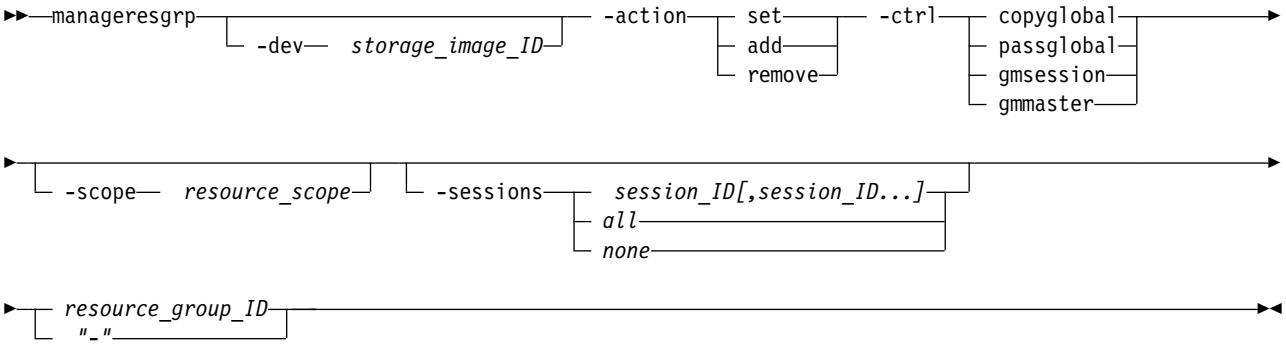
Key:

* Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

manageresgrp

The **manageresgrp** command allows you to manage the contents of any resource group object on a storage image, except resource group 0 (RG0). RG0 is predefined and cannot be created, deleted, or modified. By default, all resources belong to RG0 unless otherwise specified.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. For example, IBM.2107-75FA120. The storage image ID is required if you do not specify a fully qualified key group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-action *set* | *add* | *remove*
(Required) Specifies the action to apply to the specified resource control.

set Explicitly sets the value of a particular control to the specified value. For example, on the *copyglobal* control, it indicates that the Resource Scope specified by the **-scope** parameter replaces the existing value for the *copyglobal* control.

add Adds the specified values to the existing values for a particular control. For example, on the **gmmaster** control, it indicates that the session IDs specified by the**-sessions** parameter are added to the existing list of session IDs that are allowed to be Global Mirror Master sessions for this resource.

remove
Removes the specified values from the existing values for a particular control. For example, on the **gmmaster** control, it indicates that the session IDs specified by the **-sessions** parameter are removed from the existing list of session IDs that are allowed to be Global Mirror Master sessions for this resource.

-ctrl *copyglobal* | *passglobal* | *gmmaster* | *gmsession*
(Required) Specifies the resource group control that is the object of the specified action. These controls can be physical (for example, a volume), logical (for example, a Global Mirror relationship), or they can be behavioral (for example, the allowed Copy Services behavior between two resources).

copyglobal
Specifies the Copy Services Global Resource Scope (CS GRS). This resource scope applies to the Establish PPRC Pair (mkpprc) and Establish FlashCopy Pair (mkflash) operations that are issued from any source (a network user ID or a host system). The primary/source logical volume of the volume pair verifies that the secondary/target logical volume is within the scope of the CS GRS in its associated resource group and the secondary/target logical volume

of the volume pair verifies that the primary/source logical volume is within the scope of the CS GRS in its associated resources group. If either check fails, the requested operation is rejected.

passglobal
Specifies the Pass-Through Global Resource Scope (PGRS). Some host-issued commands can be issued to a given logical volume but operate on a different logical volume or logical subsystem that is specified in the command parameters. In this case, the logical volume that receives the command is called the pass-through device and the logical volume or the logical subsystem that the command operates on is called the destination device (logical volume or logical subsystem). This resource scope applies to any Copy Services commands from any source (network user ID or a host system) that result in a pass-through operation. In this case, the resource group label of the destination device must be within the scope of the pass-through logical volume PGRS. Pass-through occurs for the following situations:

1. A Copy Services request is issued to a CKD logical volume that operates on either a fixed block logical volume or logical subsystem. In cases where the Copy Services request establishes a pair or a path, it is the primary/source logical volume or logical subsystem that the must be within the scope of the pass-through logical volume.
2. A Copy Service request is issued to a CKD logical volume that is a PPRC primary of a PPRC pair, the request specifies that the request is issued to the PPRC secondary of the PPRC pair, and the request operates on either a fixed block logical volume or logical subsystem. In cases where the Copy Services request establishes a pair or a path, it is the primary/source logical volume or logical subsystem that the must be within the scope of the pass-through logical volume.

gmmaster
Specifies the Global Mirror Masters Allowed control, which specifies one or more session numbers whose associated global mirror master is allowed to be managed through an LSS or LCU associated with this resource group. Each global mirror session has one and only one master and it may be run on any storage image that has an LSS associated with the session. The Global Mirror Masters Allowed control allows the user to select which storage image the GM master for a given session is allowed to run on by allowing the GM master in a resource groups on one or more storage images and disallowing it in the resource groups on one or more other storage images. In order for a GM master to managed through an LSS or LCU associated with this resource group, both the Global Mirror Master Allowed and the Global Mirror Sessions Allowed controls must both allow the session number associated with the GM master.

gmsession
Specifies the Global Mirror Sessions Allowed control, which specifies one or more session numbers that are allowed to be managed through an LSS or LCU associated with this resource group. Each LSS or LCU has one at most one assigned session number. The logical volumes associated with the LCU or LSS are either associated with the session number of the LCU or LSS, or with no session number. The Global Mirror Sessions Allowed control allows the user to partition the available GM session numbers between the set of resource groups that are each associated with a given tenant in a multi-tenancy environment.

-scope *resource_scope*
(Optional) Specifies the resource scope, which must meet the following criteria:

- Must be 1 to 32 characters long.
- The characters are limited to upper and lower-case alphabetic, numeric, and the special characters dash (-), underscore (_), and period (.).

Required when the **-ctrl** specifies *copyglobal* or *passglobal*.

-sessions *session_ID[,session_ID]* | *all* | *none*
(Optional) (Required when the **-ctrl** parameter specifies *gmmasterl* or *gmsession*.) Specifies one or more Global Mirror session IDs for the specified control. A Session ID is a hexadecimal number in the

01 - FF range. To specify a range of session IDs, separate the session IDs with a hyphen (-). You must separate multiple session IDs or ranges of session IDs with a comma between each ID or range of IDs. If you specify **-sessions all**, session IDs 01-FF are used. If you specify **-sessions none**, no sessions are used.

resource_group_ID | –
(Optional) The resource group ID. The resource group ID begins with the letters *RG* and ends with a decimal number. If the resource group ID is not specified, one will be assigned. If you use the dash (-), the specified value is read from standard input.

The table below includes all of the valid combinations of the manageresgrp command and parameter combinations and their effects on the DS8000. Commands with the same control and action can be combined.

manageresgrp command	**Effects**
-ctrl copyglobal -action set -scope **(RS)**	All Copy Services tasks that establish, or re-establish, a relationship are subject to the specified resource scope **(RS)**
-ctrl gmmaster -action add -sessions **(list)**	Specifies **(list)** of Global Mirror sessions that are allowed to be a Global Mirror Master sessions for this resource. The specified list is added to any existing list for Global Mirror Master sessions.
-ctrl gmmaster -action remove -sessions **(list)**	Specifies **(list)** of Global Mirror sessions that are allowed to be a Global Mirror Master sessions for this resource. The specified list is removed from any existing list for Global Mirror Master sessions.
-ctrl gmmaster -action set -sessions **(list)**	Specifies **(list)** of Global Mirror sessions that are allowed to be a Global Mirror Master sessions for this resource. The specified list replaces any existing list for Global Mirror Master sessions.
-ctrl gmsession -action add -sessions **(list)**	Specifies **(list)** of Global Mirror sessions that are allowed to be a Global Mirror sessions for this resource. The specified list is added to any existing list for Global Mirror sessions.
-ctrl gmsession -action remove -sessions **(list)**	Specifies **(list)** of Global Mirror sessions that are allowed to be a Global Mirror sessions for this resource. The specified list is removed from any existing list for Global Mirror sessions.
-ctrl gmsession -action set -sessions **(list)**	Specifies **(list)** of Global Mirror sessions that are allowed to be a Global Mirror sessions for this resource. The specified list replaces any existing list for Global Mirror sessions.
-ctrl passglobal -action set -scope **(RS)**	All Copy Services tasks that establish, or re-establish, a relationship through a pass through device are subject to the specified resource scope **(RS)**

Notes:

1. When a Copy Services command is rejected as a result of a resource group policy attribute, the indicated error code or error message specifies which resource involved in the operation has caused the operation to be rejected. As such, the policy that cause the rejection would be in the resource group that is associated with that resource. For example, an Establish PPRC Pair (mkpprc) operation causes a PPRC to be established between a primary volume A and a secondary volume B. If the request is rejected with an error that indicates that the primary volume has rejected the operation because of a target resource scope error, then it is the resource group of volume A that contains the CS_GRS parameter that caused the operation to be rejected.
2. The combination of the CS_GRS and P_GRS controls allow the user to limit the set of resources that can be accessed by Copy Services operations issued to a given logical volume. A host (by volume

group or other host dependent mechanism) or user ID (by URS) is limited to a specific set of connection logical volumes. The CS_GRS limits what target/secondaries the connection volume is allowed to have. The P_GRS limits what destination volumes can be accessed from either the connection volume or its associated secondary, when the command is issued to the secondary through the primary. For example a FICON host can only issue commands to volumes defined in the HCD of the host. If it issues a command to logical volume A that requests volumes A and B to establish a PPRC pair, the CS GRS in the RGs of volumes A and B are checked to see whether the relationship is allowed. If the relationship is established and the host then issues a command to the PPRC secondary that request that fixed block volumes C and D establish a flash copy pair, the PGRS of volume B is checked to see whether it allows volume B to passthrough to volume C, and then subsequently the CS_GRS in the RGs of volumes C and D are checked to see whether the relationship is allowed.

Example

Invoking the manageresgrp command

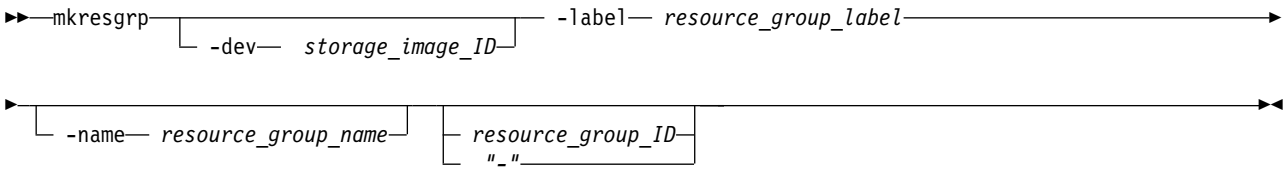
```
dscli> manageresgrp
–dev IBM.2107-75FA120 –action set –ctrl copyglobal –scope Product_A RG1
```

The resulting output

```
Resource Group RG1 successfully modified.
```

mkresgrp

The **mkresgrp** command creates a resource group object on a storage image.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. For example, IBM.2107-75FA120. The storage image ID is required if you do not specify a fully qualified resource group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-label *resource_group_label*
(Required) Specifies the resource group label. The resource group label is 1 to 32 characters and is limited to upper and lower case alphabetic and numeric characters, and the special characters (-), (_), and (.). Resource group labels must be unique within the domain of the specified device.

-name *resource_group_name*
(Optional) Specifies the user assigned nickname for this resource group object. The maximum length is 64 single-byte or 32 double-byte characters.

resource_group_ID | **-**
(Optional) The resource group ID. The resource group ID begins with the letters *RG* and ends with a decimal number. If the resource group ID is not specified, one will be assigned. If you use the dash (-), the specified value is read from standard input.

Example

Invoking the mkresgrp command

```
dscli> mkresgrp -dev IBM.2107-75FA120
-label Product_A -name "A_Group" RG1
```

The resulting output

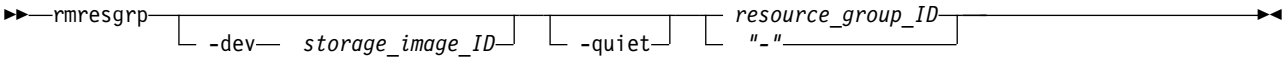
```
CMUC00133I mkresgrp: Resource Group RG1 successfully created.
```

rmresgrp

The **rmresgrp** command removes a resource group object on a storage image.

Notes:

1. Resource Group 0 (zero) is predefined and cannot be created, deleted, or modified. By default, all resources belong to this group unless otherwise specified.
2. You cannot delete a resource group if there are resources still assigned to that resource group. For example, if a volume or LCU has a resource group ID of RG4, then you cannot delete RG4 until you remove those volumes or LCUs from the resource group.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. For example, IBM.2107-75FA120. The storage image ID is required if you do not specify a fully qualified resource group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-quiet
(Optional) Turns off the resource group removal confirmation prompt.

resource_group_ID | **-**
(Required) An array of one or more resource group IDs or resource group ID ranges to be removed. A resource group ID range is defined by two resource group IDs that are separated by a hyphen. Multiple resource group IDs or resource group ID ranges must be separated with a blank space between each ID.

The resource group ID begins with the letters *RG* and ends with a decimal number. If you use the dash (-), the specified value is read from standard input.

Example

Invoking the rmresgrp command

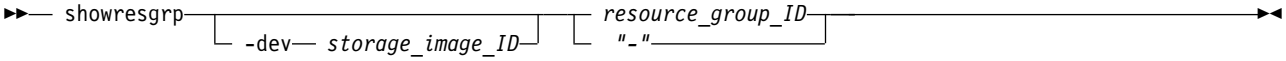
```
dscli> rmresgrp -quiet -dev IBM.2107-75FA120 RG1
```

The resulting output

```
Resource Group RG1 successfully deleted.
```

showresgrp

The **showresgrp** command displays detailed properties of a resource group.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

resource_group_ID | **-**
(Optional) The resource group ID. The resource group ID begins with the letters *RG* and ends with a decimal number. If the resource group ID is not specified, one will be assigned. If you use the dash (-), the specified value is read from standard input.

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

Invoking the showresgrp command

```
dscli> showresgrp RG1
```

The resulting output

Attribute	Value
ID	RG1
Name	Product_A.admins
State	Normal
Label	Product_A
CS_Global _RS	Product_A
Passthru_ Global_RS	PUBLIC
GM_sessions _allowed	01-FF
GM_masters _allowed	7,20-2F,A0

Report field definitions

ID Indicates the unique identifier that is assigned to this resource group ID.

name Indicates the nickname that you assigned for this resource group object.

state Indicates the current configuration state of this resource group. One of the following values is displayed:

Normal
Indicates that the resource group is not being configured.

Configuring
Indicates that the resource group is being configured.

Configuration Error
Indicates that the resource group configuration process failed to complete successfully.

Deconfiguring
Indicates that a resource group is in the process of being deleted.

Deconfiguration Error
Indicates that the resource group deletion process failed to complete successfully.

Label Indicates the resource group label. The resource group label is 1 to 32 characters and is limited to upper and lower case alphabetic and numeric characters, and the special characters (-), (_), and (.).

CS_Global_RS
Indicates all of the Copy Services requests that establish or re-establish a relationship, and are subject to this resource scope including FlashCopy and PPRC relationships.

Passthru_Global_RS
Indicates all of the Copy Services requests that are issued though a pass-through logical volume are treated as a relationship between the pass-through logical volume and source logical volume and are subject to this resource scope.

GM_Sessions_Allowed
Indicates an array of Global Mirror session IDs that are allowed to be used for the volumes in this resource.

GM_Masters_Allowed
Indicates an array of Global Mirror session IDs that are allowed to be used as a master session for volumes in this resource.

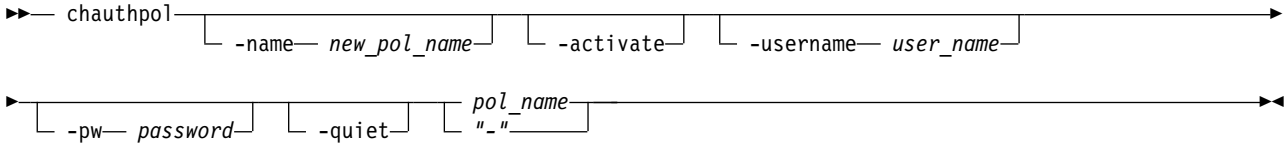
Security commands

Commands that are used to maintain security, including user accounts and security and private network security are provided for DS command-line interface.

chauthpol

The **chauthpol** command changes the general attributes of an authentication policy, such as the policy name and the activation state.

General attributes are attributes that apply to every policy type. To change specific attributes that apply only to some policy types, use the **setauthpol** command.



Parameters

-name *new_pol_name*
(Optional) Specifies the new name for the authentication policy. This parameter cannot be used with the **-activate** parameter.

-activate
(Optional) Activates the specified authentication policy. This parameter requires the **-username** and **-pw** parameters, and optionally the **-quiet** parameter. This parameter cannot be used with the **-name** parameter. To use this parameter, you must be logged in with the currently active authentication policy with storage administrator authority. The name that is specified by the **-username** parameter must have storage administrator authority in the specified authentication policy to be activated.

-username *user_name*
(Optional) Specifies a valid user name in the policy that is being activated.

Note: You must have administrator privileges in the current policy to modify the user name in the new policy.

This parameter is required when the **-activate** parameter is specified.

-pw *password*
(Optional) Specifies the password of the user name that is specified in the **-username** parameter.

Note: You must have administrator privileges in the current policy to modify the password in the new policy.

This parameter is required when the **-activate** parameter is specified.

-quiet
(Optional) Turns off the modification confirmation prompt for this command.

pol_name | –
(Required) Specifies the name of the authentication policy. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 1

Invoking the chauthpol command to change the attributes of an authentication policy.

```
dscli> chauthpol –quiet –name my_policy1 my_policy_01
```

The resulting output

```
Authentication policy my_policy_01 successfully modified.
```

Example 2

Invoking the chauthpol command to change the attributes of an authentication policy.

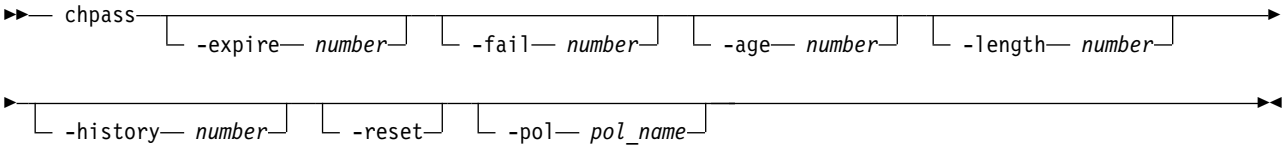
```
dscli> chauthpol –quiet
–activate –username admin –pw test2 my_policy2
```

The resulting output

```
Authentication policy my_policy_2 successfully modified.
```

chpass

The **chpass** command changes the password expiration time, minimum password age, minimum password length, password history, and the number of login attempts for a basic user account.



Parameters

-expire *number*
(Optional) Specifies the number of days a user account password is valid before it expires and needs to be changed by the user. The default number of days is 90. If you do not want the password to expire, enter a value of zero. After the password expires, the user is not allowed to log in unless the password is changed.

Note: Though the expiration range is 0 - 9999, a nonzero must also be greater than or equal to the minimum age value. So, to avoid a possible password problem, do not allow the password to expire before you change the value. For example, assuming an age value of 7 (meaning you cannot change the password more than once a week), but the expiration is set to 5 days, there would be two days in which you could not use the account.

-fail *number*
(Optional) Specifies the number of login attempts allowed on any given user account. The number of login attempts can be zero to twenty-five. The default number of login attempts is 5. If you do not want a limit on the number of login attempts, enter zero. After the number of login attempts is exceeded, the user account is locked.

-age *number*
(Optional) Specifies the number of days that a user must wait before changing a password. A zero disables the age requirement of the password, and the default age is 1.

Note: The range is 0 - *expiration value*, inclusive. The real check is that the age value should not exceed the password expiration value, for the reason explained under the -expire parameter.

-length *number*
(Optional) Specifies the minimum length (from 6 to 64 characters) for passwords. The default is 8.

-history *number*
(Optional) Specifies the number of unique passwords (from 1 to 16) that a user must go through before reusing a password. The default is 8.

-reset
(Optional) Resets all of the password rule values to their default values.

Note: The **-reset** parameter can only be specified by itself and not with any other parameters.

-pol *pol_name*
(Optional) Specifies the name of the basic authentication policy. This parameter is optional if you have authenticated with a 'basic' authentication policy type, but it is required if you are authenticated with another type of authentication policy.

Example

Invoking the chpass command

```
dscli> chpass -pol my_policy1
-expire 20 -fail 0 -age 0 -length 8 -history 10
```

The resulting output

```
Password parameters with my_policy1 successfully changed.
```

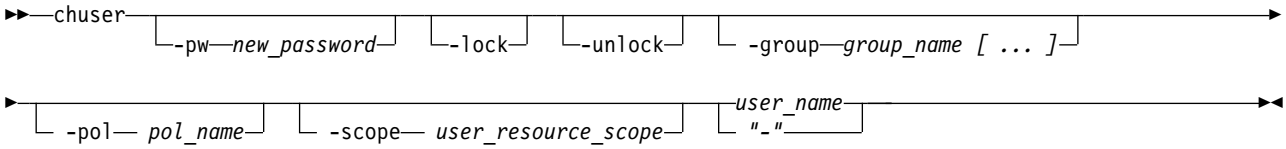
chuser

The **chuser** command modifies and locks or unlocks a DS CLI or a DS8000 Storage Management GUI basic user account.

All users can use this command to change their own password. However, users with administrator authority can use this command to update any user's account password to modify user group authority, or lock or unlock a user account. Users with administrator authority can also modify passwords and user groups that are based on IBM Engineering and IBM Service roles.

Users with security administrator authority can only modify the attributes of a user ID in the security administrator group role. Users with administrator authority can modify the attributes of a user ID in any user role except security administrator.

Note: When a person with administrator authority designates the password, the password is set to expire upon its initial use. The user of the password must use the **chuser** command to establish a new password before access to the rest of the DS CLI can be granted.



Parameters

-pw *new_password*
(Optional) Specifies that the new designated password is to be assigned to the user.

Notes:

1. When a person with administrator authority uses this parameter in association with the **-unlock** parameter, the new password is temporary and expires upon the initial use.
2. When a person without administrator authority uses this parameter, the new password becomes the valid password and replaces the prior password.

new_password
Indicates the new password.

The new password must meet the following criteria:

- Must be at least the minimum length as set by an administrator and no longer than 64 characters.
- Must contain at least two types of characters from the three groups: alphabetic, numeric, and symbols.
  - Allowable characters include a-z, A-Z, 0-9, and the symbols!@#$%&*().
- Cannot contain the user ID of the user.

Note: If symbols are contained in your password, you might be required to enclose the password in quotation marks. This prevents any special interpretations or expansions by the operating system shell program.

Note: Even with a valid password, a user cannot interactively log in when all of the following conditions are present:

- The version of DS CLI used is pre-R6.1
- Entering the password without either the -passwd or -pwfile parameters
- The DS CLI is operating in the Windows (all versions) or NetWare environments
- The password contains anything other than alphabetic or numeric characters (that is, symbols)

But if any of these conditions are not present, then the user should not encounter any problems in logging in with a valid password.

-lock
(Optional) Specifies that a user account is to be locked.

A person with administrator authority can use this parameter to lock a user account. The locking action occurs when the user authenticates the account. If a user is already active (authenticated) and is using the DS CLI, the lock does not occur until logout.

-unlock
(Optional) Specifies that a user account is to be unlocked.

A person with administrator authority can use this parameter to unlock a user account when the user can no longer log in to the DS CLI. A person might not be able to log in to the DS CLI for the following reasons:

- The user forgot the password and in an attempt to log in went beyond the set number of allowable attempts. Going beyond the set limit locks the user account.

  Note: When unlocking a user account for this scenario, the administrator must also assign a new password to the user using the **-pw** parameter. The new password is temporary and immediately expires after its initial use. The administrator must notify the user of this new password.
- Someone with administrator authority has locked the user account.

-group *group_name [...]*
(Optional) Specifies a user's access authority group or groups. A user can be assigned to many of the following user groups:
- admin (Administrator)
- ibm_engineering (Engineer)
- op_storage (Physical Operator)
- op_volume (Logical Operator)
- op_copy_services (Copy Services Operator)
- secadmin (Security Administrator)
- ibm_service (Service)
- monitor (Monitor)
- no_access (No Access)

-pol *pol_name*
(Optional) Specifies the name of the basic authentication policy. This parameter is optional if you have authenticated with a 'basic' authentication policy type, but it is required if you are authenticated with another type of authentication policy.

-scope *user_resource_scope*
(Optional) Specifies the user resource scope, which must meet the following criteria:
- Must be 1 to 32 characters long
- The characters are limited to upper and lower case alphabetic, numeric, and the special characters, dash (-), underscore (_), and period (.). You can also define the scope as a single asterisk (*).

user_name | **-**
(Required) Specifies the name of the user account.

Notes:
1. The administrator inserts the name of the user account that is affected by the changes (that is, group name, lock, or unlocking).
2. Users who are changing their passwords can insert their user names.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Examples

Invoking the chuser command with password parameter

```
dscli> chuser -pw xy0abcde testuser
```

The resulting output

```
User testuser successfully modified.
```

Invoking the chuser command with group parameter

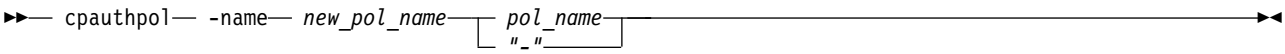
```
dscli> chuser -group ibm_service cuscus
```

The resulting output

```
CMUC00134I User cuscus successfully modified.
```

cpauthpol

The **cpauthpol** command copies an existing authentication policy to a new policy. You can only copy a policy within the same storage complex. Copying between different storage complexes is not supported.



Parameters

-name *new_pol_name*
(Required) Specifies the new name for the (copied) authentication policy.

pol_name | **-**
(Required) Specifies the name of the existing authentication policy that is being copied. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the cpauthpol command to copy an existing authentication policy to a new policy.

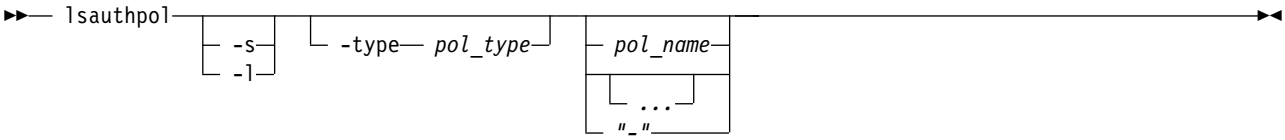
```
dscli> cpauthpol –name my_policy4 my_policy2
```

The resulting output

```
New authentication policy my_policy4 successfully copied.
```

lsauthpol

The **lsauthpol** command displays a list of all the authentication policies on the storage complex.



Parameters

-s
(Optional) Displays only the policy name. You cannot use the **-s** and the **-l** parameters together.

-l
(Optional) Displays the default output and location. You cannot use the **-s** and the **-l** parameters together.

-type *pol_type*
(Optional) Specifies the type of authentication policy that you want to list. For example, SAS.

SAS The storage authentication service (SAS) authentication type uses the authentication service that is provided by the IBM System Storage Productivity Center.

Basic The Basic authentication type uses the DS8000 Basic authentication service.

pol_name ... | –
(Optional) Specifies the name of the authentication policy that you want to list. The ellipsis (...)

indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the lsauthpol command to display a list of all the authentication policies on the storage complex.

```
dscli> lsauthpol -l
```

The resulting output

```
Name       Type  State    Location
Pol1       SAS   inactive http://yoursite.com
Pol2       SAS   active   http://yoursite.com
my_policy1 Basic inactive 9.12.133.155
```

Report field definitions

Name*
Indicates the name of the authentication policy.

Type
Indicates the authentication policy type. (For example, SAS.)

State
Indicates the state of the authentication policy, either active or inactive.

Location+
Indicates the URL or the IP address for the authentication server. Multiple locations are separated by commas without spaces.

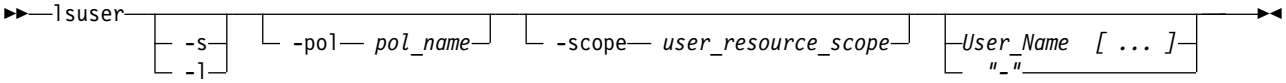
Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

lsuser

The **lsuser** command returns a list of basic user account names and access authority levels.



Parameters

-s
(Optional) Displays only the user names. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output. You cannot use the **-l** and the **-s** parameters together.

-pol *pol_name*
(Optional) Specifies the name of the basic authentication policy. This parameter is optional if you authenticated with a 'basic' authentication policy type, but it is required if you are authenticated with another type of authentication policy.

-scope *user_resource_scope*
(Optional) Displays only the user accounts that have the specified user resource scope. If you do not specify the scope, users with any user resource scope might be displayed.

User_Name ... | -
(Optional) Displays the user accounts with the specified user names. Separate multiple user names with a blank space between each name.

Example

Note: For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as table.

Invoking the lsuser command

```
dscli> lsuser -l –pol my_policy1
```

The resulting output

Name	Group	State
admin	admin	active
engineering	ibm_engineering	active
inge	ibm_engineering	active
monitor	monitor	active
nick	op_volume,op_copy_services	active
noaccess	no_access	active
opvol	op_volume	active
physicalop	op_storage	active
robmel	admin	active
service	ibm_service	active

Report field definitions

Name*
Indicates the user name that is assigned to the user account.

Group
Indicates the access authority group of the user. One or more of the following group designations is displayed:

- admin (Administrator)
- secadmin (Security Administrator)
- ibm_engineering (Engineer)
- op_storage (Physical Operator)
- op_volume (Logical Operator)
- op_copy_services (Copy Services Operator)
- ibm_service (Service)
- monitor (Monitor)
- no_access (No Access)

State
Indicates the status of the user account for the designated user group, either active or locked.

Scope+
Indicates the user resource scope. A dash (-) is displayed if the storage unit does not support resource groups.

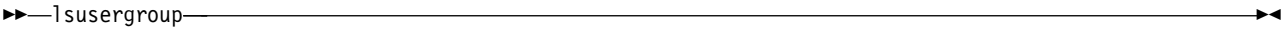
Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

lsusergroup

The **lsusergroup** command displays user groups and permissions by user roles.



Parameters

This command has no parameters.

Examples

```
dscli> lsusergroup
```

The resulting output

```
Name
============
123321
admin
ibm_engineering
ibm_service
monitor
no_access
op_copy_services
op_storage
op_volume
secadmin
```

Report field definitions

Name
Indicates the name of the user group.

Roles
Indicates user roles and the permissions by user role.

admin (Administrator)
This user role (also referred to as the storage administrator) has the highest level of authority. It allows a user to add or remove user accounts. This role has access to all service functions and DS8000 resources, and all DS8000 storage image resources except those functions that are reserved for Security Administrator users. Users in this role cannot be assigned to any other role, and users in any other role might not be assigned to this role.

The administrator role can name and define custom roles by selecting individual tasks that are permitted to each custom role. For example, the output of this command includes "123321," which is an example of a custom role created by the administrator.

ibm_engineering
This user role is typically assigned to IBM support personnel who perform all service functions and other functions that might be needed. This role does not have access to the logical configuration or data on the storage system.

ibm_service
This user role is typically assigned to IBM support personnel who service the hardware (install, remote, or repair) and update firmware. This role does not have access to the logical configuration or data on the storage system. Users in this role cannot be assigned to any other role, and users in any other role cannot be assigned to this role.

monitor
This user role has access to all read-only, non-security service functions and all DS8000 resources. Users in this role cannot be assigned to any other role, and users in any other role cannot be assigned to this role.

no_access
This user role has no access to any service functions or DS8000 resources. This role is the default selection and is assigned to a user account that is not associated with any other user role. Users in this role cannot be assigned to any other role, and users in any other role cannot be assigned to this role.

op_copy_services (Copy Services operator)
This user role has access to all Copy Services functions and resources, excluding security functions. This role might be assigned in combination with the Logical Operator role, but not in combination with any other role.

op_storage (Physical operator)
This user role has access to resources that are related to physical configuration, including storage complex, storage unit, storage image, management console, arrays, ranks, and extent pools. The Physical Operator role does not have access to security functions. Users in this role cannot be assigned to any other role, and users in any other role cannot be assigned to this role.

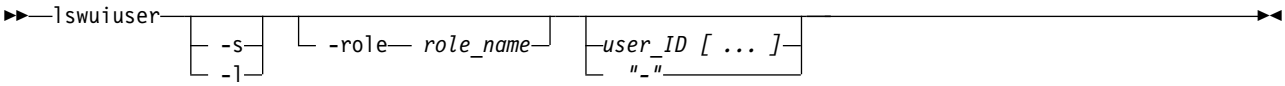
op_volume (Logical operator)
This user role has access to resources that relate to logical volumes, hosts, host ports, logical subsystems, and volume groups, excluding security functions. This role can be assigned in combination with the Copy Services Operator role, but not in combination with any other role.

op_volume + op_copy_services (Logical and copy operator)
This user role includes the permissions for both the logical operator and copy operator roles. This role has access to resources that relate to logical volumes, hosts, host ports, logical subsystems, and volume groups, excluding security functions. This role also has access to all Copy Services functions and resources, excluding security functions.

secadmin (Security administrator)
This user role allows users to initiate recovery key operations, and add other users to this role. Users in this role cannot be assigned to any other role, and users in any other role cannot be assigned to this role.

lswuiuser

The **lswuiuser** command displays users who have access to the DS Service GUI on the HMC.



Parameters

-s
(Optional) Displays only the user IDs. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output. You cannot use the **-l** and the **-s** parameters together.

-role *role_name*
(Optional) Displays only users with access to the DS Service GUI on the HMC, with a specific task role. If you do not specify a task role, users with access to the DS Service GUI, with any task role, might be displayed.

user_ID ... | -
(Optional) Displays one or more user IDs of the users with access to the DS Service GUI on the HMC. Specify multiple user IDs with a blank space between each name.

Example

Note: For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as table.

Invoking the lswuiuser command

```
dscli> lswuiuser -l
```

The resulting output

User ID	Task Role	Authentication Type	Description
PE	ibm_engineering	Local authentication	
CE	ibm_service	Local authentication	HMC user
Customer		Local authentication	Customer

Report field definitions

User ID*
Indicates the user ID of the user on the DS Service GUI on the HMC.

If you specify the -s parameter, then this is the only column that displays.

Task Role
Indicates the task roles of the users that have access to the DS Service GUI on the HMC. This column displays with or without the **-l** parameter specified.

Authentication Type
Indicates the authentication type of the users that have access to the DS Service GUI on the HMC. This column displays with or without the **-l** parameter specified.

Description
Indicates the description that is associated with the user that has access to the DS Service GUI on the HMC. This column displays if you specify the **-l** parameter.

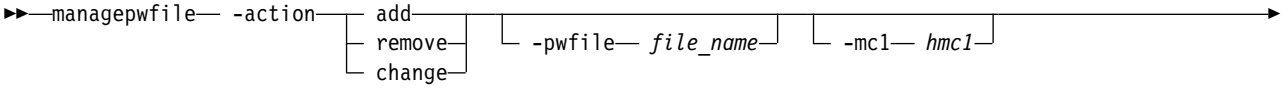
Key:

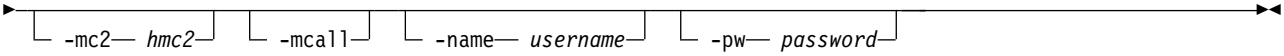
* Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

managepwfile

The **managepwfile** command manages a security file that contains passwords for any existing DS user account. This password security file can then be used with the **-pwfile** parameter of the **dscli** command instead of the **-passwd** parameter. Although using a security file is not required, it is strongly recommended as a more secure method of entering the user's password when invoking the **dscli** command. The security file method does not require entering the password on the command line, nor is it contained in any script files.





Parameters

-action
(Required) Specifies that a process that is designated by the sub-parameter be enacted on the password file.

add Specifies that an entry for the user and the specified primary HMC be created in the password file. If the password file does not exist, it will be created.

remove
Specifies that the password file entry be removed for the designated user.

change
Specifies that the password file entry be changed for the designated user.

-pwfile *file_name*
(Optional) Specifies the name that you want to use for the password file. You can specify the password file as an absolute path or a relative path. The relative path is obtained from the current working directory.

-mc1 *hmc1*
(Optional) Specifies the DNS name or the IP address. *hmc1* designates the Model 2107 primary HMC, Model 1750 primary SMC, Model 2105 primary Copy Services server DNS or IP address.

Note: If you do not specify this parameter, the DS CLI will use the value that was specified for **-hmc1** in the current CLI session connection, or the default value, if specified, for HMC1 in your profile file. This value, as entered and converted to lower case, along with the value of the **-name** parameter is used as a key in the password file. If the values for **-mc1** and **-mc2** are equivalent, then only one key is generated.

-mc2 *hmc2*
(Optional) Specifies the DNS name or the IP address of the secondary HMC.

Note: If you do not specify this parameter, the DS CLI will use the value that was specified for **-hmc2** in the current CLI session connection, or the default value, if specified, for HMC2 in your profile file. This value, as entered and converted to lower case, along with the value of the **-name** parameter is used as a key in the password file. If the values for **-mc1** and **-mc2** are equivalent, then only one key is generated.

-mcall
(Optional) Specifies any existing DNS name or IP address in the password file. The specified action is applied to any entry in the password file that is identified only by the **–name** value as the lookup key. You can use this parameter to change all of the passwords for a particular user in one command.

For the remove and change actions, you can specify this value instead of the **–mc1** parameter, the **-mc2** parameter, or both. For the add action, this parameter is ignored.

-name *username*
(Optional) Specifies the name that you use to access the DS CLI. This information, along with the **-mc1** and **-mc2** parameter information, is used as keys in the password file.

-pw *password*
(Optional) Specifies a user-assigned password. The password is case-sensitive.

Notes:

1. The primary or secondary Storage Manager IP addresses, along with the user name, are used to form a key to locate the user's password in the security file. However, these values are stored as character

strings. Therefore, either use the same string every time for the IP addresses, or use **managepwfile** to store all variations of the IP addresses. In other words, the IP address of `9.11.64.211` and the DNS name of `myds.chicago.abc.com` would form different keys even though they identify the same machine.

2. A password file is created with a user's default protection mask. The user can update the protection mask to allow access only to the owner of the file. Also, you must write down the directory name where the password file is contained in case you need to use it later.
3. The password file has a default value of *<user_home>*/dscli/security.dat.

   The home directory *<user_home>* is defined by the Java system property named "user.home" The location of your password file is determined by your operating system. The following examples are home directories in different operating systems:

   Windows XP operating system
   For a Windows XP operating system, the property value defaults to the environment variable **%USERPROFILE%**. As a result, your personal profile is `C:\Documents and Settings\<username>\dscli\security.dat`.

   UNIX or Linux operating system
   For an UNIX or Linux operating system, the property value defaults to the environment variable **$HOME**. As a result, your personal profile is `~/dscli/security.dat`.

   i5/OS
   For the i5/OS, your personal profile is `/home/<username>/dscli/security.dat`.

   Note: The values of the Java system properties can be redefined by JRE options. If you are having problems, check to see whether you have an environment setting like the following on your local system:

   ```
   _JAVA_OPTIONS=-Duser.home=...
   ```
4. In some circumstances this command might return more than one error/informational message.

Example

Invoking the managepwfile command

```
dscli> managepwfile -action add -mc1 myess.ibm.com  -mc2 myess2.ibm.com
-name testuser –pw AB9cdefg
```

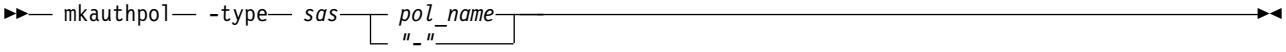
The resulting output

```
Record myess.ibm.com/testuser successfully added to password file
c:\Documents and Settings\testuser\dscli\security.dat
Record myess2.ibm.com/testuser successfully added to password file
c:\Documents and Settings\testuser\dscli\security.dat
```

mkauthpol

The **mkauthpol** command creates an empty authentication policy.

After you create an authentication policy, use the **manageauthpol** command to configure it, and the **chauthpol** command to enable it.



Parameters

-type *sas*
(Required) Specifies the authentication policy type. Currently, SAS (Storage Authentication Service) is the only valid value for this parameter and it is required.

pol_name | -
(Required) Specifies the name of the authentication policy. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the mkauthpol command to an empty authentication policy.

```
dscli> mkauthpol -type sas my_policyname
```

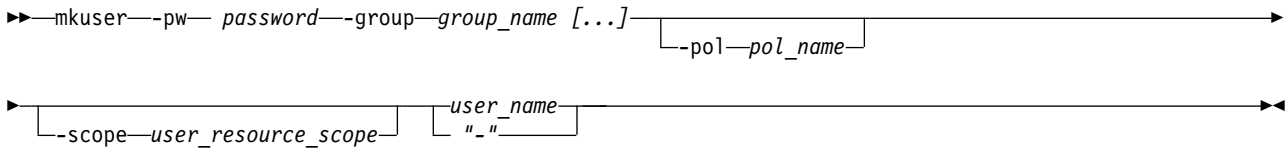
The resulting output

```
Authentication policy my_policyname successfully created.
```

mkuser

The **mkuser** command creates a DS CLI or a DS Storage Manager Basic authentication policy user account with a password and user group authority.

Only users with security administrator authority can create user IDs in the security administrator role and cannot create an ID in any other group role. A user with administrator authority can create user IDs in any group role except security administrator. A user with administrator authority can also create users with Service and Engineering roles, which are typically assigned to IBM support personnel.



Parameters

-pw *password*
(Required) Specifies the password that is assigned to the user that allows their use of the DS CLI command-line function. This password is temporary and expires after the initial use. The user must assign themselves a new password using the **chuser** command before they can use any other commands in the DS CLI.

password
The password that is assigned by the administrator to a user.

The password must meet the following criteria:

- Must be at least the minimum length as set by an administrator and no longer than 16 characters.
- Must contain at least two types of characters from the three groups: alphabetic, numeric, and symbols.
- Allowable characters are: a-z, A-Z, 0-9, and the symbols !@#$%&*().
- Cannot contain the user ID of the user.

Note: If symbols are contained in your password, you might be required to enclose the password in quotation marks to prevent any special interpretations or expansions by the operating system shell program.

Note: Even with a valid password, a user cannot interactively login when all of the following conditions are present:

- The version of DS CLI used is pre-R6.1

- Entering the password without either the -passwd or -pwfile parameters
- The DS CLI is operating in the Windows (all versions) or Netware environments
- The password contains anything other than alphabetic or numeric characters (that is, symbols)

But if any of these conditions is not present, then the user should not encounter any problems in logging in with a valid password.

-group *group_name [...]*
(Required) Specifies the user's access authority group.

- admin (Administrator)
- ibm_engineering (Engineer)
- op_storage (Physical Operator)
- op_volume (Logical Operator)
- op_copy_services (Copy Services Operator)
- secadmin (Security Administrator)
- ibm_service (Service)
- monitor (Monitor)
- no_access (No Access)

Note: An ellipsis [...] signifies that multiple comma-separated values can be specified.

-pol *pol_name*
(Optional) Specifies the name of the basic authentication policy. This parameter is optional if you have authenticated with a 'basic' authentication policy type, but it is required if you are authenticated with another type of authentication policy.

-scope *user_resource_scope*
(Optional) Specifies the user resource scope, which must meet the following criteria:

- Must be 1 to 32 characters long
- The characters are limited to upper and lower case alphabetic, numeric, and the special characters, dash (-), underscore (_), and period (.). You can also define the scope as a single asterisk (*).

The default scope is * for users in both the administrator and the security administrator groups, and *PUBLIC* for users in all other authority groups.

Example: Product_A

Note: The user resource scope is matched to one or more resource group IDs that are assigned to resource groups. If the resource group ID of a resource group matches the user resource scope, the user is authorized to issue Copy Services requests to a logical volume, LSS, or LCU that is assigned to the resource group. To issue a Copy Services request to establish a volume pairing, an LSS-pairing, or LCU-pairing, the user must be authorized to access the source volume, source LSS, or source LCU, respectively. To issue a Copy Services request that operates on an LSS or LCU or has a session parameter, the user must be authorized to access that LSS or LCU.

user_name | **-**
(Required) The user name of the new user that you are creating.

An account name must be between 1 and 64 characters in length and can include the alphanumeric characters (a-z, A-Z, 0-9), the period(.), the hyphen (-), the underscore (_), and letters from any other languages. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Examples

Invoking the mkuser command

```
dscli> mkuser -pw AB9cdefg
-group service,op_copy_services -pol my_policy1 testuser
```

The resulting output

```
CMUC00133I mkuser: User testuser with my_policy1
successfully created.
```

Invoking the mkuser command

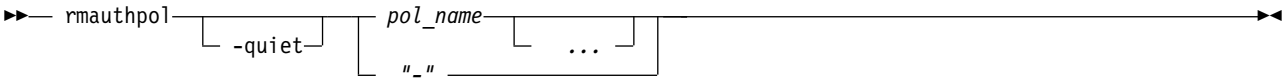
```
dscli> mkuser -pw ********* -group ibm_engineering ingeniero
```

The resulting output

```
CMUC00133I mkuser: User ingeniero successfully created.
```

rmauthpol

The **rmauthpol** command allows you to remove an authentication policy.



Parameters

-quiet
(Optional) Turns off the removal confirmation prompt for this command.

pol_name ... | -
(Required) Specifies the name of the authentication policy that you want to remove. The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmauthpol command to remove an authentication policy.

```
dscli> rmauthpol my_policy3
```

The resulting output

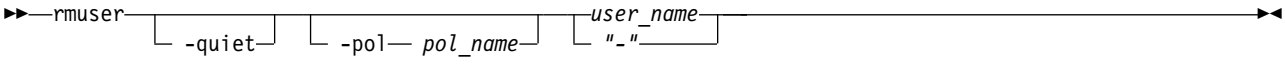
```
Are you sure you want to delete Authentication policy my_policy3? y/n y

Authentication policy my_policy3 successfully removed.
```

rmuser

The **rmuser** command removes a basic user account.

Only users with administrator authority can use this command to delete a user ID. Users with only security administrator authority can delete a user ID in the security administrator group role. Users with administrator authority can delete a user ID in any group role except security administrator.



Parameters

-quiet
(Optional) Turns off the removal confirmation prompt for this command.

-pol *pol_name*
(Optional) Specifies the name of the basic authentication policy. This parameter is optional if you authenticated with a 'basic' authentication policy type, but it is required if you are authenticated with another type of authentication policy.

user_name | –
(Required) Specifies the user name of the user account to be removed.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmuser command

```
dscli> rmuser –pol my_policy1 testuser
```

The resulting output

```
Are you sure you want to delete User_Name testuser? y/n
Y
User_Name testuser within my_policy1 successfully deleted.
```

Invoking the rmuser command

```
dscli> rmuser engineering
```

The resulting output

```
Are you sure you want to delete user engineering? y/n
Y
User engineering successfully deleted.
```

Invoking the rmuser command

```
dscli> rmuser service
```

The resulting output

```
Are you sure you want to delete user service? y/n
Y
User service successfully deleted.
```

Invoking the rmuser command

```
dscli> rmuser ingeniero
```

The resulting output

```
Are you sure you want to delete user ingeniero? y/n
Y
User ingeniero successfully deleted.
```

setauthpol

The **setauthpol** command modifies policy attributes that apply to a specific type of authentication policy, changing the contents of the policy. To change attributes that are independent of the policy type, use the **chauthpol** command.

With the **setauthpol** command, you can map external users and groups to one or more authority groups and to a user resource scope. Only users with security administrator authority can map a user ID or

group to the security administrator group role. Users with security administrator authority cannot be mapped to any other group role. A user with administrator authority can map a user ID, or group, to any group role except security administrator.

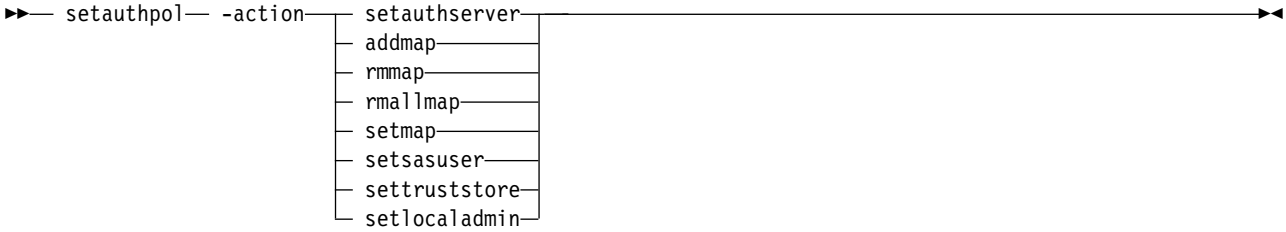
Note: The **setauthpol** command cannot be run from the Embedded DS CLI window.

Using the **setauthpol** command, storage administrators with global resource scope authority can enable or disable a local administrator for the authentication policy. The local administrator is a specified user account in the local security repository who can log in to the system when a given remote user directory policy is configured and the remote user directory server is not accessible.

Depending on the policy type and the action that is selected, all of the other parameters can vary in meaning. For this reason, the syntax diagrams and the parameter descriptions are separated by policy types and actions. If a parameter is not found under a specific policy type, then it does not apply to that policy type.

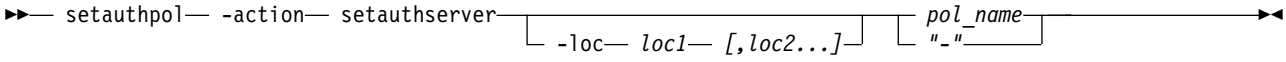
Notes:

- You must have administrator or security administrator authority and a user resource scope of '*' in the current policy to use this command.
- If an external user belongs to several external groups that map to more than one user resource scope, other than DEFAULT, the user cannot log on unless there is also a mapping between the external user and one specific user resource scope.
- The previous parameters -groupmap and -usermap, used with the addmap, rmmap, and setmap actions, are now deprecated but are still valid for use in commands. The new parameters -extgroup, -extuser, and -dsgroup replace the deprecated parameters, and cannot be used in the same command line with them.

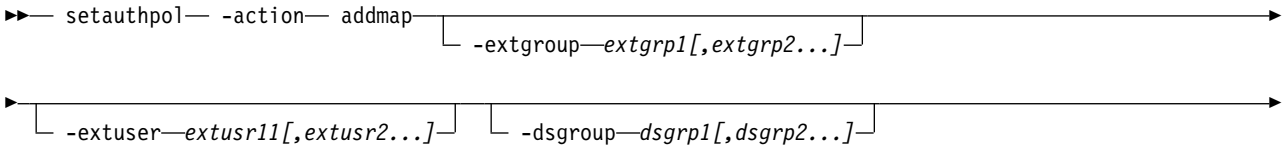


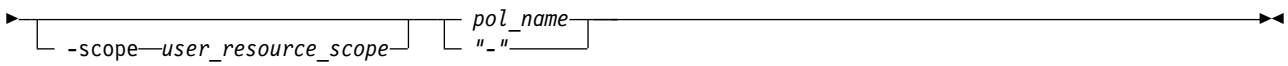
Each of the following sections shows the options for one of the listed **-action** parameters:

-action setauthserver

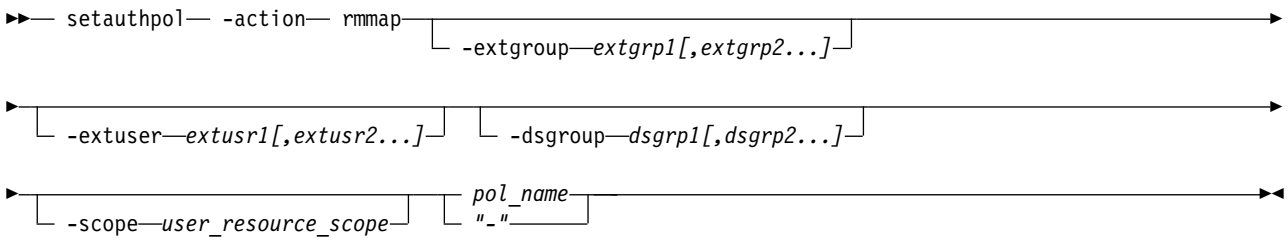


-action addmap

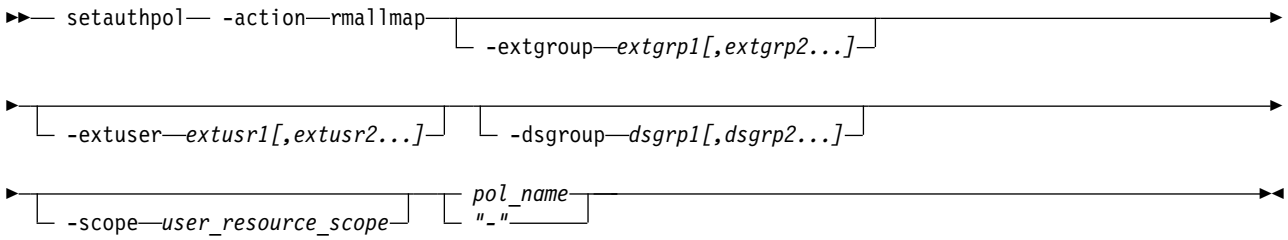




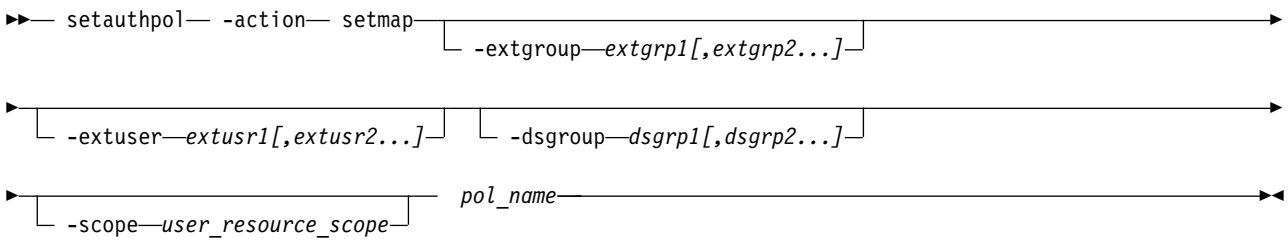
-action rmmap



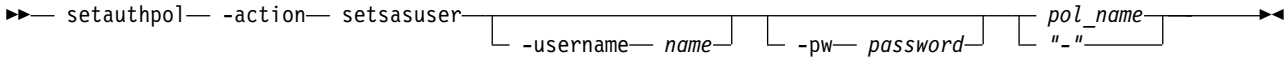
-action rmallmap



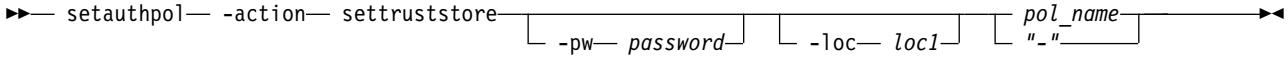
-action setmap



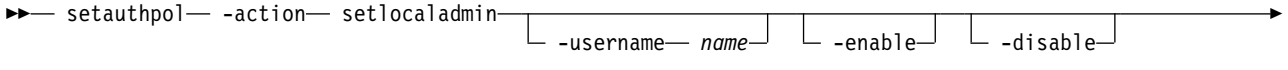
-action setsasuser

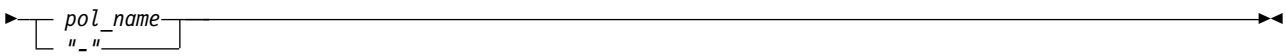


-action settruststore



-action setlocaladmin





Action parameters

-action setauthserver
(Required) Specifies the authentication server that is used in the policy.

-action addmap
(Required) Adds to the mappings of external users or groups in the DS8000 authority group roles.

Note: Either the **-extgroup** or **-extuser** (or both) parameters can be used with either the **-dsgroup** or **-scope** (or both) parameters to add a mapping between the specified mapping pairs.

-action rmmap
(Required) Removes either all or specific mappings of external users or groups from the DS8000 authority group roles.

Note: Either the **-extgroup** or **-extuser** (or both) parameters can be used with either the **-dsgroup** or **-scope** (or both) parameters to remove the mapping between the specified mapping pairs. Any existing mappings that are not part of the specified parameters are not removed.

-action rmallmap
(Required) Removes multiple mappings of either DS8000 authority group roles and user resource scope to external users or groups.

Note: Any, or all, of the **-extgroup**, **-extuser**, **-dsgroup**, or **-scope** parameters can be used to remove all mapping pairs with any of the specified parameters. Any existing mappings that are not part of the specified parameters are not removed.

-action setmap
(Required) Maps external users or groups to DS8000 authority group roles. If previous mappings are defined, the *setmap* action replaces them. Use the *addmap* action to add new mappings without replacing previous versions. All unspecified roles are unchanged.

Note: Either the **-extgroup** or **-extuser** (or both) parameters can be used with either the **-dsgroup** or **-scope** (or both) parameters to set the mapping between the specified mapping pairs. Any existing mappings for the specified keyword pairs are replaced by the specified mapping.

-action setsasuser
(Required) Specifies the storage authentication service (SAS) user.

-action settruststore
(Required) Specifies the location of the truststore file. You must specify the password to check the integrity of the keystore data, and to set the truststore.

-action setlocaladmin
(Required) Specifies that a storage administrator with global resource scope authority can set a local administrator for the authentication policy. The local administrator is specified as the user account in the local security repository who can log in to the system when a given remote user directory policy is configured and the remote user directory server is not accessible.

The following table includes all the valid combinations of actions and parameters and their effects on the DS8000 system. One or more options from parameter group 1 must be specified. If any options are listed in parameter group 2, one or more of those options must also be specified.

Table 12. setauthpol action and parameter combinations

Action	Parameter group 1	Parameter group 2	Effects
addmap	-extgroup *extgrp1*[,*extgrp2*] -extuser *extusr1*[,*extusr2*]	-dsgroup *dsgrp1*[,*dsgrp2*] -scope *user_resource_scope*	Add mapping to existing maps
rmmap	-extgroup *extgrp1*[,*extgrp2*] -extuser *extusr1*[,*extusr2*]	-dsgroup *dsgrp1*[,*dsgrp2*] -scope *user_resource_scope*	Remove mapping from existing maps
rmallmap	-extgroup *extgrp1*[,*extgrp2*] -extuser *extusr1*[,*extusr2*] -dsgroup *dsgrp1*[,*dsgrp2*] -scope *user_resource_scope*		Remove specified values from all existing maps
setauthserver	-loc *loc1*[,*loc2*]		Set location of authentication server
setmap	-extgroup *extgrp1*[,*extgrp2*] -extuser *extusr1*[,*extusr2*]	-dsgroup *dsgrp1*[,*dsgrp2*] -scope *user_resource_scope*	Specify mapping to replace existing maps
setsasuser	-username *name*	-pw *password*	Set SAS user name and password
settruststore	-loc *loc1*	-pw *password*	Set location and password of trust store file
setlocaladmin	-username *name*	-enable or -disable	Set the local administrator account of the specified authentication policy to enable or disable.

Parameters for a basic policy type, sorted by the selected action

No attributes that can be set for this policy type are available.

Parameters for a SAS policy type, sorted by the selected action

Parameters for **-action** setauthserver:

-loc *loc1[,loc2...]*
(Optional) Specifies the URL location of the authentication servers. *loc1* and *loc2* are URLs specified as an IPv4, IPv6, or DNS-named IP address. Multiple locations are separated by commas without spaces. SAS servers use the HTTPS protocol and port number 16311, both of which are specified in the URL. For example, `https://9.11.236.10:16311/TokenService/services/Trust`

pol_name | **-**
(Required) Specifies the name of the authentication policy. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Parameters for **-action** setmap:

-extgroup *extgrp1[,extgrp2...]*
(Optional) This action maps DS8000 authentication group roles (**-dsgroup**) and/or a user resource scope (**-scope**) to a specified list of external authentication group names (**extgrp**). Multiple group names are separated with commas without spaces. All unspecified external groups are unchanged.

Example: ESS_ds1,ESS_ds2

-extuser *extusr1[,extusr2...]*
(Optional) This action maps DS8000 authentication group roles (**-dsgroup**) and/or a user resource scope (**-scope**) to the specified list of external authentication user names (**-extuser**). Multiple user names are separated with commas without spaces. All unspecified external users are unchanged.

Example: fred,sally

-dsgroup *dsgrp1[,dsgrp2...]*
(Optional) This action lists DS8000 authentication group roles (**dsgrp**) that consist of one or more of the following role names: "admin", "secadmin", "op_storage", "op_volume", "op_copy_services", "monitor", "no_access" "ibm_engineering", and "ibm_service". Multiple role names are separated with commas without spaces.

Note: The "ibm_engineering" and "ibm_service" roles cannot be combined with other existing roles. The DSCLI allows users to have multiple roles, but currently the only combination allowed is "op_volume" + "op_copy_services."

Example: op_volume,op_copy_services

-scope *user_resource_scope*
(Optional) Specifies the user resource scope, which must meet the following criteria:

- Must be 1 - 32 characters long
- The characters are limited to upper and lowercase, alphabetic, numeric, and the special characters, dash (-), underscore (_), and period (.). You can also define the scope as a single asterisk (*).

The default scope is * for users in the administrator and security administrator authority groups, and *PUBLIC* for users in all other authority groups.

Example: Product_A

Note: The user resource scope is matched to one or more resource group IDs that are assigned to resource groups. If the resource group ID of a resource group matches the user resource scope, the user is authorized to issue Copy Services requests to a logical volume, LSS, or LCU that is assigned to the resource group. To issue a Copy Services request to establish a volume pairing, an LSS-pairing, or LCU-pairing, you must be authorized to access the source volume, source LSS, or source LCU, respectively. To issue a Copy Services request that operates on an LSS or LCU or has a session parameter, you must be authorized to access that LSS or LCU.

pol_name | **-**
(Required) Specifies the name of the authentication policy. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Parameters for **-action** addmap:

-extgroup *extgrp1[,extgrp2...]*
(Optional) This action maps DS8000 authentication group roles (**-dsgroup**) and user resource scope (**-scope**) to a specified list of external authentication group names (**extgrp**). Multiple group names are separated with commas without spaces. All unspecified external groups are unchanged.

Example: ESS_ds1,ESS_ds2

-extuser *extusr1[,extusr2...]*
(Optional) This action maps DS8000 authentication group roles (**-dsgroup**) and user resource scope (**-scope**) to the specified list of external user (**-extuser**). Multiple user names are separated with commas without spaces. All unspecified external users are unchanged.

Note: Use of this parameter is not recommended for maintenance reasons.

Example: fred,sally

-dsgroup *dsgrp1[,dsgrp2...]*
(Optional) This action lists DS8000 authentication group roles (**dsgrp**) that consist of one or more of the following role names: "admin", "secadmin", "op_storage", "op_volume", "op_copy_services", "service", "monitor", and "no_access". Multiple role names are separated with commas without spaces.

Example: op_volume,op_copy_services

-scope *user_resource_scope*
(Optional) Specifies the user resource scope, which must meet the following criteria:

- Must be 1 - 32 characters long
- The characters are limited to upper and lower case alphabetic, numeric, and the special characters, dash (-), underscore (_), and period (.). You can also define the scope as a single asterisk (*).

The default scope is * for users in the administrator authority group, and *PUBLIC* for users in all other authority groups.

Example: Product_A

Note: The user resource scope is matched to one or more resource group IDs that are assigned to resource groups. If the resource group ID of a resource group matches the user resource scope, the user is authorized to issue Copy Services requests to a logical volume, LSS, or LCU that is assigned to the resource group. To issue a Copy Services request to establish a volume pairing, an LSS-pairing, or LCU-pairing, you must be authorized to access the source volume, source LSS, or source LCU, respectively. To issue a Copy Services request that operates on an LSS or LCU or has a session parameter, you must be authorized to access that LSS or LCU.

pol_name | **-**
(Required) Specifies the name of the authentication policy. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Parameters for **-action** rmmap:

-extgroup *extgrp1[,extgrp2...]*
(Optional) This action unmaps DS8000 authentication group roles (**-dsgroup**) and user resource scope (**-scope**) to a specified list of external authentication group names (**extgrp**). Multiple group names are separated with commas without spaces. All unspecified external groups are unchanged.

Example: ESS_ds1,ESS_ds2

-extuser *extusr1[,extusr2...]*
(Optional) This action unmaps DS8000 authentication group roles (**-dsgroup**) and user resource scope (**-scope**) to the specified list of external users (**-extuser**). Multiple user names are separated with commas without spaces. All unspecified external users are unchanged.

Note: Use of this parameter is not recommended for maintenance reasons.

Example: fred,sally

-dsgroup *dsgrp1[,dsgrp2...]*
(Optional) This action lists DS8000 authentication group roles (**dsgrp**) that consist of one or more of the following role names: "admin", "secadmin", "op_storage", "op_volume", "op_copy_services", "service", "monitor", and "no_access". Multiple role names are separated with commas without spaces.

Example: op_volume,op_copy_services

-scope *user_resource_scope*
(Optional) Specifies the user resource scope, which must meet the following criteria:

- Must be 1 - 32 characters long
- The characters are limited to upper and lowercase alphabetic, numeric, and the special characters, dash (-), underscore (_), and period (.). You can also define the scope as a single asterisk (*).

The default scope is * for users in the administrator authority group, and *PUBLIC* for users in all other authority groups.

Example: Product_A

Notes:

1. The user resource scope is matched to one or more resource group IDs that are assigned to resource groups. If the resource group ID of a resource group matches the user resource scope, you are authorized to issue Copy Services requests to a logical volume, LSS, or LCU that is assigned to the resource group. To issue a Copy Services request to establish a volume pairing, an LSS-pairing, or LCU-pairing, you must be authorized to access the source volume, source LSS, or source LCU, respectively. To issue a Copy Services request that operates on an LSS or LCU or has a session parameter, you must be authorized to access that LSS or LCU.
2. When a scope mapping is removed from a **-extuser** or **-extgroup**, the default scope will still apply.

pol_name | **-**
(Required) Specifies the name of the authentication policy. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Parameters for **-action** rmallmap:

-extgroup *extgrp1[,extgrp2...]*
(Optional) Specifies that all maps with the external groups specified by this parameter are removed. Multiple group names are separated with commas without spaces. All unspecified external groups are unchanged.

Example: ESS_ds1,ESS_ds2

-extuser *extusr1[,extusr2...]*
(Optional) Specifies that all maps with the external users specified by this parameter are removed. Multiple user names are separated with commas without spaces. All unspecified external users are unchanged.

Note: Use of this parameter is not recommended for maintenance reasons.

Example: fred,sally

-dsgroup *dsgrp1[,dsgrp2...]*
(Optional) Specifies that all maps with the DS8000 groups specified by this parameter are removed. DS8000 authentication group roles (**dsgrp**) consist of one or more of the following role names: "admin", "secadmin", "op_storage", "op_volume", "op_copy_services", "service", "monitor", and "no_access". Multiple role names are separated with commas without spaces.

Example: op_volume,op_copy_services

-scope *user_resource_scope*
(Optional) Specifies that all maps with the scope specified by this parameter are removed. The user resource scope must meet the following criteria:

- Must be 1 - 32 characters long
- The characters are limited to upper and lower case alphabetic, numeric, and the special characters, dash (-), underscore (_), and period (.). You can also define the scope as a single asterisk (*).

The default scope is * for users in the administrator authority group, and *PUBLIC* for users in all other authority groups.

Example: Product_A

Notes:

1. The user resource scope is matched to one or more resource group IDs that are assigned to resource groups. If the resource group ID of a resource group matches the user resource scope, you are authorized to issue Copy Services requests to a logical volume, LSS, or LCU that is assigned to the resource group. To issue a Copy Services request to establish a volume pairing, an LSS-pairing, or LCU-pairing, you must be authorized to access the source volume, source LSS, or

source LCU, respectively. To issue a Copy Services request that operates on an LSS or LCU or has a session parameter, you must be authorized to access that LSS or LCU.

2. When a scope mapping is removed from a **-extuser** or **-extgroup**, the default scope still applies.

pol_name | -
(Required) Specifies the name of the authentication policy. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Parameters for **-action** settruststore:

-pw *password*
(Optional) Specifies the truststore password.

-loc *loc1*
(Optional) Specifies the local truststore file location. Only one truststore location can be specified. *loc1* is the full path name of the file that is stored on the local system.

Example: c:\mystore\trust.dat

pol_name | -
(Required) Specifies the name of the authentication policy. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Parameters for **-action** *setsasuser*:

-username *name*
(Optional) Specifies the user name that is used internally by SAS (Storage Authentication Service). Only one user name can be specified.

-pw *password*
(Optional) Specifies the user name password that is used internally by SAS.

pol_name | -
(Required) Specifies the name of the authentication policy. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Parameters for **-action** *setlocaladmin*:

-username *name*
(Optional) A user name in the local security repository. Only a storage administrator with global resource scope can be specified as the local administrator for a remote authentication policy.

-enable
(Optional) Specifies to enable the local administrator of the authentication policy.

-disable
(Optional) Specifies to disable the local administrator of the authentication policy.

The -disable parameter is not valid when specified with the -enable or -username parameters.

pol_name | -
(Required) Specifies the name of the authentication policy. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

The difference between the previous and new syntax for specifying users and groups

- **Previous syntax**

```
dscli> setauthpol –action addmap –groupmap
admin:Admins/monitor:Admins,Users
```

- **New syntax**

```
dscli> setauthpol –action addmap –extgroup Admins
–dsgroup admin
```

```
dscli> setauthpol –action addmap –extgroup
Admins,Users –dsgroup monitor
```

Remove all mappings between the external group 'Dept54' and all internal DS0000 groups and/or scope

```
dscli> setauthpol –action rmallmap –extgroup Dept54
```

Starting the `setauthpol` command to modify the contents of the policy

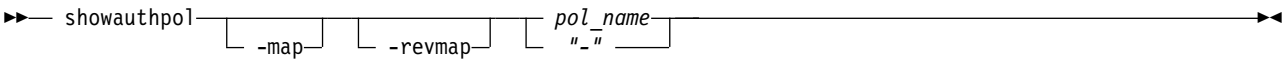
```
dscli> setauthpol –setsasuser –username was_user
–pw test2 my_policy1
```

The resulting output

```
Authentication policy my_policy1 successfully modified.
```

showauthpol

The **`showauthpol`** command displays detailed properties of a specified authentication policy.



Parameters

`-map`
(Optional) Displays tables with mappings of Basic authorization group roles and user resource scopes to external groups and users in the specified policy. No table is displayed if there are no mapping relationships in the specified policy.

Note: The **`-map`** and **`-revmap`** parameters cannot be used together.

`-revmap`
(Optional) Displays tables with mappings of external groups and users to Basic authorization group roles and user resource scopes in the specified policy. No table is displayed if there are no mapping relationships in the specified policy.

Note: The **`-map`** and **`-revmap`** parameters cannot be used together.

pol_name | **-**
(Required) Specifies the name of the authentication policy that you would like to view. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 1

Invoking the showauthpol command to view the detailed properties of a specified authentication policy.

```
dscli> showauthpol my_policy2
```

The resulting output

name	my_policy2

type	SAS
state	inactive
location	9.11.xxx.xxx
truststore	my_policy2_trustStore.jks
sasuser	-
localAdmin	admin

Example 2

Invoking the showauthpol command to view the detailed properties of a specified authentication policy using the -map parameter.

```
dscli> showauthpol -map my_policy2
```

The resulting output

name	my_policy2
type	SAS
state	inactive
location	9.11.xxx.xxx
truststore	my_policy2_trustStore.jks
sasuser	-
localAdmin	admin

```
============================Role Group Maps ============================
```

DS_group	Ext_group
op_volume	grpa,grpb
op_copy_services	grpa,grpb

```
============================Role User Maps============================
```

DS_group	Ext_user
admin	joe,bob
no_access	bob

```
============================Scope Group Maps============================
```

Scope	Ext_group
Accounting	grpa
Purchasing	grpb

```
============================Scope User Maps============================
```

Scope	Ext_user
*	bob
Accounting	joe

Example 3

Invoking the showauthpol command to view the detailed properties of a specified authentication policy using the -revmap parameter.

```
dscli> showauthpol -revmap my_policy2
```

The resulting output

name	my_policy2
type	SAS
state	inactive
location	9.11.xxx.xxx
truststore	my_policy2_trustStore.jks
sasuser	-
localAdmin	admin

```
============================Role Group Maps============================
```

Ext_group	DS_group
grpb	op_volume,op_copy_services
grpa	op_volume,op_copy_services

```
============================Role User Maps============================
```

Ext_user	DS_group
bob	admin,no_access
joe	admin

```
============================Scope Group Maps============================
```

Ext_group	Scope
grpb	Purchasing
grpa	Accounting

```
============================Scope User Maps============================
```

Ext_user	Scope
bob	*
joe	Accounting

Report field definitions

For a basic policy type, the following properties are displayed:

name
Indicates the name of the authentication policy.

type
Indicates the authentication policy type.

state
Indicates the state of the authentication policy (active or inactive).

location
Indicates the names or IP addresses of the Hardware Management Consoles that were used when users logged in. If users logged in from more than one location, a list of locations is displayed, separated by commas.

expire
Indicates the number of days a user account password is valid before it expires.

age
Indicates the minimum days a user must wait before changing a password.

fail
Indicates the number of login attempts allowed on any given user account.

length
Indicates the minimum length of a password.

history
Indicates the number of unique passwords that a user must go through before reusing a password.

For a SAS type policy, the following properties are displayed:

name
Indicates the name of the authentication policy.

type
Indicates the authentication policy type.

state
Indicates the state of the authentication policy (active or inactive).

location
Indicates the URL for the authentication server. Multiple locations are separated by commas.

truststore
Indicates the truststore file name.

sasuser
Indicates the user name used internally by SAS (Storage Authentication Service).

localAdmin
Indicates a user name that is used as the local administrator.

Note: A dash (-) means that the local administrator was not available.

The **-map** and **-revmap** are mutually exclusive, but both display the mapping from external users and groups to DS8000 user role groups and user resource scopes. The **-map** parameter displays this information from the DS8000 point of view and is useful for answering questions like, "Which external users and groups map to the DS8000 role group administrator?" The **-revmap** parameter displays the same information, but from the external point of view and is useful answering questions like, "Which DS8000 user role groups and user resource scope map to the external group Human_Resources?"

The following additional properties are displayed when the **-map** parameter is specified:

Role Group Maps

DS_group
Displays the name of the DS8000 authority group. The user authority group can consist of one or more of one of the following roles: admin, secadmin, op_storage, op_volume, op_copy_services, service, monitor, or no_access.

Ext_group
Displays the external groups that are mapped to each selected DS8000 authority group. Multiple external group names are separated by commas.

Role User Maps

DS_group
Displays the name of the DS8000 authority group. The user authority group can consist of one or more of the following roles: admin, op_storage, op_volume, op_copy_services, service, monitor, or no_access.

Ext_user
Displays the external users that are mapped to each selected DS8000 authority group. Multiple external user names are separated by commas.

Scope Group Maps

Scope
Displays the user resource scope.

Ext_group
Displays the external group names that are mapped to each user resource scope. Multiple external group names are separated by commas.

Scope User Maps

Scope
Displays the user resource scope.

Ext_user
Displays the external users that are mapped to each user resource scope. Multiple external user names are separated by commas.

The following additional properties are displayed when the **-revmap** parameter is specified:

Role Group Maps

Ext_group
Display one or more external group names.

DS_group
Display the DS8000 authority group names that are mapped to each external group name. Multiple external group names are separated by commas.

Role User Maps

Ext_user
Displays one or more external users.

DS_group
Displays the DS8000 authority group names that are mapped to each external group name. Multiple external group names are separated by commas.

Scope Group Maps

Ext_group
Displays the name of the DS8000 authority group. The user authority group can consist of one or more of the following roles: admin, op_storage, op_volume, op_copy_services, service, monitor, or no_access.

Scope
Displays the user resource scope that is mapped to each external group name.

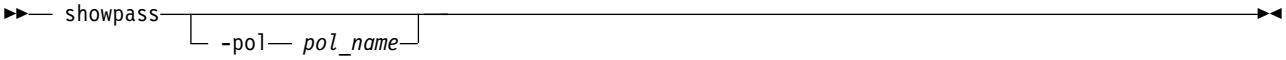
Scope User Maps

Ext_group
Displays the name of the DS8000 authority group. The user authority group can consist of one or more of the following roles: admin, op_storage, op_volume, op_copy_services, service, monitor, or no_access.

Scope
Displays the user resource scope that is mapped to each external group name.

showpass

The **showpass** command lists the properties of passwords.



Parameters

-pol *pol_name*
(Optional) Specifies the name of the basic authentication policy. This parameter is optional if you have authenticated with a 'basic' authentication policy type, but it is required if you are authenticated with another type of authentication policy.

Example

Invoking the showpass command

```
dscli> showpass -pol my_policy1
```

The resulting output

```
dscli> showpass
Password Expiration   90 days
Failed Logins Allowed 5
Password Age          1 days
Minimum Length        8
Password History      8
```

Report field definitions

Password Expiration
The number of days all user account passwords are valid before they expire.

Note: In the **showuser** command, the DaysToExpire field displays the number of days a particular user account password will be valid before it expires.

Failed Logins Allowed
The number of login attempts allowed on any given user account.

Password Age
The minimum days a user must wait before changing a password.

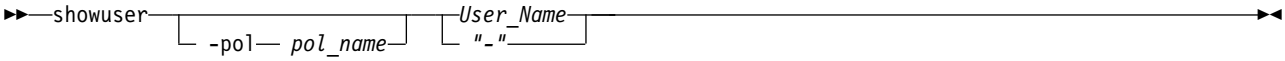
Minimum Length
The minimum length of a password

Password History
The number of unique passwords that a user must go through before reusing a password.

showuser

The **showuser** command displays details for basic user accounts.

A CLI user who has administrative authority uses this command to display the properties (group assignment, user account status and number of failed logins) that are associated with a current user account.



Parameters

-pol *pol_name*
(Optional) Specifies the name of the basic authentication policy. This parameter is optional if you have authenticated with a 'basic' authentication policy type, but it is required if you are authenticated with another type of authentication policy.

User_Name | **-**
(Required) Specifies the name of the user account .

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI show commands, the results are shown in table format for clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output reports that are associated with the **showuser** command.

Invoking the showuser command

```
dscli> showuser -pol my_policy1  testuser
```

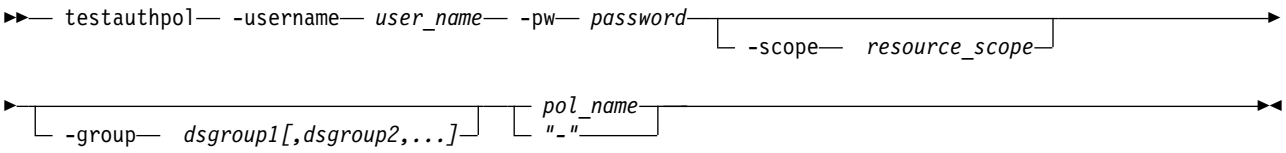
The resulting output

Column Header	Description
Name	Name of the user that you have queried.
Group	The user's access authority group. One or more of the following group designations is displayed: • admin (Administrator) • secadmin (Security Administrator) • op_storage (Physical Operator) • op_volume (Logical Operator) • op_copy_services (Copy Services Operator) • service (Service Operator) • monitor (Monitor) • no_access (No Access)
State	The status of the user account for the specified user group, either active or locked.
FailedLogin	Count of login failures since the last successful login for this user. This number resets to 0 with each successful login.

Column Header	Description
DaysToExpire	The number of days a user account password is valid before it expires. **Note:** A dash (-) means the field is unsupported on the server. 9999 means the password expiration function is disabled.
Scope	The user resource scope. A dash (-) is displayed if the storage unit does not support resource groups.

testauthpol

The **testauthpol** command allows you to test the specified authentication policy.



Parameters

-username *user_name*
(Required) Specifies the user name for the authentication policy that is being tested. For example, if the current policy is Policy1 and you want to test Policy2, then you must be logged in with an administrator user account in Policy1, and provide a valid user name and password for Policy2.

-pw *password*
(Required) Specifies the password for the user name in the policy being tested.

-scope *resource_scope*
(Optional) Specifies the expected scope that the user is associated with. The test will succeed if the user is associated with the scope. The scope mappings can be set and changed using the **setauthpol** command.

-group *dsgroup1[,dsgroup2,...]*
(Optional) Specifies the expected groups that the user belongs to. The test will succeed if the user is part of each of the specified groups. The group mappings can be set and changed using the **setauthpol** command.

pol_name | –
(Required) Specifies the authentication policy that you want to test. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the testauthpol command to test a specified authentication policy.

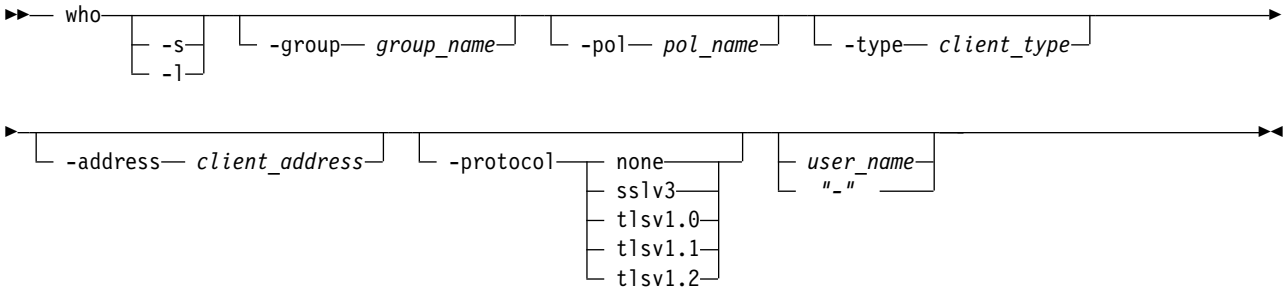
```
dscli> testauthpol –username admin –pw test2 my_policy2
```

The resulting output

```
Authentication policy my_policy2 successfully verified.
```

who

The **who** command displays authentication information for the users who are currently logged in. It also includes information about users who are logged in with service and engineering roles.



Parameters

-s
(Optional) Displays the user names for the users who are currently logged on. The **-l** and **-s** parameters cannot be used together.

-l
(Optional) Displays details about the users who are currently logged on, including user name and authority groups that the user belongs to. The **-l** and **-s** parameters cannot be used together.

-group *group_name*
(Optional) Displays the list of users who are currently logged in and who are part of the specified access authority group. If the user has multiple group roles, the user is displayed if any of those roles match the specified group.

group_name
The following list provides the list choices that can be assigned to a user. The group_name can be one of the following roles:

- admin (Administrator)
- ibm_engineering (Engineer)
- ibm_service (Service)

  Note: The "ibm_engineering" and "ibm_service" roles cannot be combined with other existing roles. The DSCLI allows users to have multiple roles, but currently the only combination allowed is "op_volume" + "op_copy_services."
- op_storage (Physical Operator)
- op_volume (Logical Operator)
- op_copy_services (Copy Services Operator)
- secadmin (Security Administrator)
- service (Service Operator)
- monitor (Monitor)
- no_access (No Access)

-pol *pol_name*
(Optional) Displays the list of users who are currently logged in under the specified client type.

-type *client_type*
(Optional) Displays the list of users who are currently logged in and who have the specified client type. One of the following client types are displayed:

DSCLI
DS Command Line Interface

DSGUI
DS Storage Management GUI

HMTU
IBM Easy Tier Heat Map Transfer Utility

SC IBM Spectrum Control™

CSM IBM Copy Services Manager

Unknown
Unknown might be displayed for older versions of any of the client types that are listed above.

-address *client_address*
(Optional) Displays the users who are currently logged in with the specified client address.

-protocol none |sslv3 | tlsv1.0 |tlsv1.1 | tlsv1.2
(Optional) Displays the users who are currently logged in with the specified client connection protocol.

None The connection is not encrypted.

sslv3 Secure Socket Layer version 3.

tlsv1.0 Transport Layer Security version 1.0.

tlsv1.1 Transport Layer Security version 1.1.

tlsv1.2 Transport Layer Security version 1.2.

user_name | **-**
(Optional) Displays only the users names for the specified user account.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the who command to view authentication information for the users who are currently logged in.

```
dscli> who
```

The resulting output

Name	Group	Policy	Scope	Type
Testuser	service,op_copy_services	localPolicy	ABC_Co	DSGUI
Biguser	admin	localPolicy	*	DSCLI
Shyuser	ibm_service	initalPolicy	Public	DSCLI

Version	Address	Protocol
-	9.123.236.89	Internal
7.6.18.204	9.123.456.139	TLSv1.2
7.8.98.409	9.123.456.139	TLSv1.2

Report field definitions

Name Indicates the user names for the users who are currently logged on.

Group Indicates the authority groups that the current users are assigned to.

Policy Indicates the name of the authentication policy that was active when each of the current users authenticated. The currently active policy can differ.

Scope Indicates the user resource scope.

Type Indicates the client type. The value is *Unknown* when the DS CLI is pre-R6.1.

Version
Indicates the client version. The value is *Unknown* when the DS CLI is pre-R6.1.

Address
Indicates the client address.

Protocol
The type of client connection protocol is provided.

None
The connection is not encrypted.

Internal
The connection is wholly contained within the HMC.

SSLv3
Secure Socket Layer version 3.

TLSv1.0
Transport Layer Security version 1.0.

TLSv1.1
Transport Layer Security version 1.1.

TLSv1.2
Transport Layer Security version 1.2.

- The dash (-) indicates that the connection protocol is unknown or the reporting query is not supported.

whoami

The **whoami** command displays authentication information for the current user, including groups with users who are logged in with service and engineering roles.



Parameters

-s
(Optional) Provides the user name of the current user. The **-l** and **-s** parameters cannot be used together.

-l
(Optional) Provides details about the current user, including user name and authority groups that the user belongs to. The **-l** and **-s** parameters cannot be used together.

Example

Invoking the whoami command to view authentication information for the current user.

```
dscli> whoami
```

The resulting output

Name	Group	Policy	Scope
admin	admin	localPolicy	*
shyguy	ibm_service	initialPolicy	PUBLIC

Report field definitions

Name Indicates the user name for the current user.

Group Indicates the authority groups that the current user is assigned to.

Policy Indicates the name of the authentication policy that was active when the current user authenticated. The currently active policy can differ.

Scope Indicates the user resource scope.

Data encryption and security commands

Use data encryption and security commands to configure the DS8000 system.

For security purposes, encryption keys are stored on external key servers, not on the DS8000 system. The **mkkeymgr**, **chkeymgr**, **rmkeymgr**, and **lskeymgr** commands are used to specify the location of the external key servers and which servers are to be used by the DS8000 system. If multiple servers are specified, it is assumed that the servers themselves manage the process to ensure that the stored keys are synchronized. Because multiple manufacturers' products might be using the same key servers, the **mkkeygrp**, **rmkeygrp**, **lskeygrp**, and **showkeygrp** commands are used to specify a *label* for any specific encryption key.

In some environments, there might be two disjointed groups of external key servers that are defined and that cannot synchronize their stored keys securely. In this case, you can specify a second *label*, one label for each group of servers. Under certain unusual circumstances, losing access to the encrypted data on the DS8000 system might be possible. This loss of access might occur if all of the external keys servers go down, or if all physical connections are lost between the DS8000 system and the external key servers. To prevent any of these possibilities from becoming a permanent loss of data access, you are required to create an encryption data access recovery key that is managed with a dual control process described in the "User account and security commands" section. The encryption recovery key itself is manually managed with the **managereckey**, **mkreckey**, and **rmreckey** commands.

The following data encryption and security commands are available:

chkeymgr
Updates the attributes of the key server entry on the storage complex.

lskeygrp
Displays a list of the key server encryption key group entries on the specified storage image.

lskeymgr
Displays a list of the key server entries that are on the storage complex.

managekeygrp
Allows you to manage an encryption key group.

managekeymgr
Allows you to manage an existing encryption key server.

managereckey
Allows you to manage an existing encryption recovery key.

mkkeygrp
Creates an entry for the key server encryption key group on the storage image.

mkkeymgr
Creates an entry for the key server on the storage complex.

mkreckey
Allows you to create an encryption recovery key.

rmkeygrp
Removes an entry for the key server encryption key group on a specified storage image.

rmkeymgr
Removes a key server entry on the storage complex.

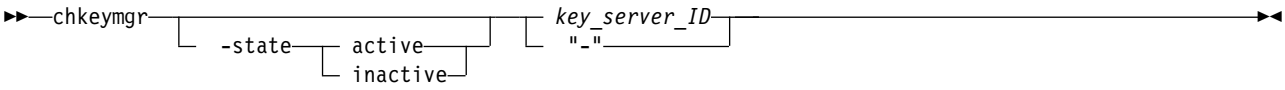
rmreckey
Allows you to remove an encryption recovery key.

showkeygrp
Displays detailed information for a specified key server encryption key group entry on the storage image.

showkeymgr
Displays detailed properties of a specified key server entry.

chkeymgr

The **chkeymgr** command updates the attributes of the key server entry on the storage complex.



Parameters

-state active | inactive
(Optional) Updates the state of the key server.

key_server_ID | **-**
(Required) Specifies the key server ID. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Invoking the chkeymgr command

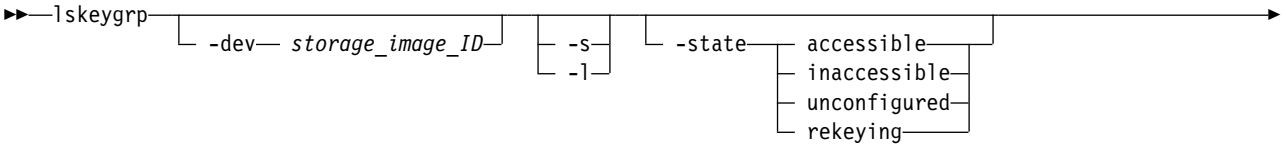
```
dscli> chkeymgr -state inactive 1
```

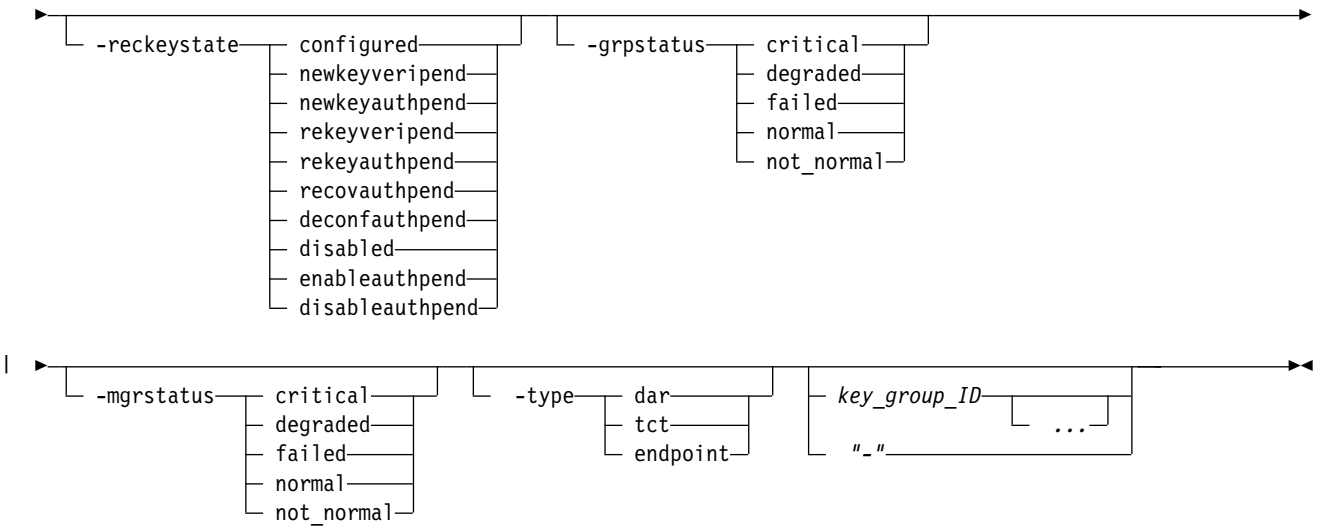
The resulting output

```
The key server 1 configuration has been changed.
```

lskeygrp

The **lskeygrp** command displays a list of the key server key group entries on a specified storage image.





Parameters

-dev *storage_image_ID*
Displays the storage image ID, which includes manufacturer, machine type, and serial number. For example, IBM.2107-75FA120. The storage image ID is required if you do not specify a fully qualified key group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the storage system is aware of more than one storage image.

Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-s
(Optional). Displays only the attributes that are identified as short output. You cannot use the **-s** and the **-l** parameters together.

-l
(Optional). Displays the default output and extra attributes that are identified as long output. You cannot use the **-l** and the **-s** parameters together.

-state accessible | inaccessible | unconfigured | rekeying
(Optional) Displays the state of the key group.

-reckeystate configured | newkeyveripend | newkeyauthpend | rekeyveripend | rekeyauthpend | recovauthpend | deconfauthpend | disabled | enableauthpend | disableauthpend
(Optional) Displays the key groups with the specified recovery key state.

-grpstatus critical | degraded | failed | normal | not_normal
(Optional) Displays the key access status for all key servers that are associated with the specified key group. The value `not_ normal` displays for all key groups whose key access status is not "normal" or the state is inactive.

-mgrstatus critical | degraded | failed | normal | not_normal
(Optional) Displays the key server path access status for all key servers that are associated with the specified key group. The value `not_ normal` displays for all key groups whose key server path summary status is not "normal" or the state is inactive.

-type dar | tct | endpoint
(Optional) Displays the type of encryption that is used by the key group:

dar Encryption for data at rest, which encrypts data that is stored on your storage system. This value is the default.

tct Encryption for transparent cloud tiering, which ensures that data is encrypted before it is

transferred to cloud storage. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

endpoint
Encryption for IBM Fibre Channel Endpoint Security, which establishes authenticated communication and encryption of data in flight for Fibre Channel connections between an IBM Z host and the storage system. The connections are secured by Fibre Channel security protocols and key server authentication that uses communication certificates. If both the host and storage system use a connection with Fibre Channel ports that support encryption, the connection will transmit encrypted data between the ports.

key_group_ID ... | -
(Optional) Displays the ID for the key group that you want to view. The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

```
dscli> lskeygrp
```

The resulting output

```
ID   state        reckeystate   reckeydate    datakeydate
=========================================================
1    accessible   configured    04/28/2015    03/18/2018

keyprotocol  type   name
============================
 KMIP        DAR    dar_encr
```

Invoking the lskeygrp command to display key server information.

```
dscli> lskeygrp -dev IBM.2107-75FA120 -l 0001
```

The resulting output

```
ID   state        reckeystate   reckeydate    datakeydate
=========================================================
1    accessible   configured    04/28/2015    03/18/2018


grpstatus    mgrstatus   label        label2         keyprotocol
=================================================================================
critical     normal      CompanyABC   CompanyABC2    KMIP
```

Report field definitions

ID Indicates the key group ID.

state
Indicates one of the following states of the key group:

accessible
The key group is accessible if it is configured and the storage image has the key group from the key server for the specified key group.

inaccessible
The key group is inaccessible if the storage image was unable to obtain the key group from the key server.

unconfigured
The key group is unconfigured if it has not been configured.

rekeying
The key group is accessible and rekeying if it is configured and the storage image has the encryption key from the key server for the key group and is in the middle of rekeying.

reckeystate
Indicates one of the following states of the recovery key:

configured
A new recovery key was requested, verified, and authorized.

unconfigured
A recovery key was not created.

newkeyveripend
A new recovery key was requested but not verified.

newkeyauthpend
A new recovery key was requested and verified, but not authorized.

rekeyveripend
A new recovery key action was requested but not verified.

rekeyauthpend
A new recovery key action was requested and verified, but not authorized.

recovauthpend
A recover action was requested, but not authorized.

deconfauthpend
A deconfigure action was requested, but not authorized.

disabled
A recovery key was disabled, and the key group is used without a recovery key.

enableauthpend
An enable action was requested, but not authorized.

disableauthpend
A disable action was requested, but not authorized.

reckeydate
Indicates the date of the last recovery key creation.

datakeydate
Indicates the date of the last data key creation. If the key group is unconfigured, then any displayed date is to be considered erroneous data.

grpstatus
Indicates one of the following values of the key group access status:

Note: The **grpstatus** parameter is refreshed by either a background process that runs every 8 hours or after the **managekeygrp** command with –action testaccess processes. The following values that display depend on either process to determine when the **grpstatus** parameter is refreshed.

For example, if all key servers were recovered from an abnormal status, the **grpstatus** parameter displays `failed` because the status reflects either process that last completed unless a background process started or the **2** command with –action testaccess has processed again.

In addition, when a key group is created, the status always displays a "-" to indicate that the background process has not started or the **managekeygrp** command with –action testaccess has not processed.

critical
The key group has access to the key on a single key server and it represents a potential single point of failure. Use the **showkeygrp** command with the –access parameter to determine the access status for each key server on the HMCs.

degraded
The key group has access to the encryption key on two or more key servers, but not all key servers. Use the **showkeygrp** command with the –access parameter to determine the access status for each key server.

failed
The key group does not have access to the encryption key on any key server. Use the **showkeygrp** command with the –access parameter to determine the access status for each key server on the HMCs.

normal
The key group has access to the encryption key on all key servers.

"-"
The dash (-) indicates that the key group state is either unconfigured, rekeying, or was created but the background process has not started or the **managekeygrp** command with –action testaccess has not processed.

mgrstatus
Indicates one of the following values of the key server path access status:

critical
At least one key server for this key group reported an access status of `normal`, `degraded`, or `critical`.. Use the **showkeymgr** command with the –access parameter to determine the access status of each HMC.

Note: A DS8000 system with only one HMC configured displays status as `normal`.

degraded
At least two key servers for the specified key group reported an access status of `normal` or `degraded`. Use the **showkeygrp** command with the –access parameter to determine the access status of each HMC.

failed
All key servers for this key group reported an access status of failed. Use the **showkeymgr** command with the –access parameter to determine the access status of each HMC.

normal
All key servers for this key group reported an access status of normal.

"-"
The dash (-) indicates that the state of the key group is either unconfigured or all key servers report a state of inactive.

label
Indicates the label for the key server key group. Because of the possible length of the label value, this column is the second to last column even as new columns are added to the output. Example MyCompany

label2
Indicates the second label for the key server key group. Because of the possible length of the label2 value, this column is the last column even as new columns are added to the output. Example MyCompany2

keyprotocol
Indicates the key server protocol that the storage system communicates with to provide key server management operations.

IPP
The storage system communicates with the specified key server using the IBM Proprietary Protocol (IPP).

KMIP
The storage system communicates with the specified key server using the Key Management Interoperability Protocol (KMIP).

type
Indicates the type of encryption that is used by the key group:

DAR Encryption for data at rest, which encrypts data that is stored on your storage system. This value is the default.

TCT Encryption for transparent cloud tiering, which ensures that data is encrypted before it is transferred to cloud storage. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

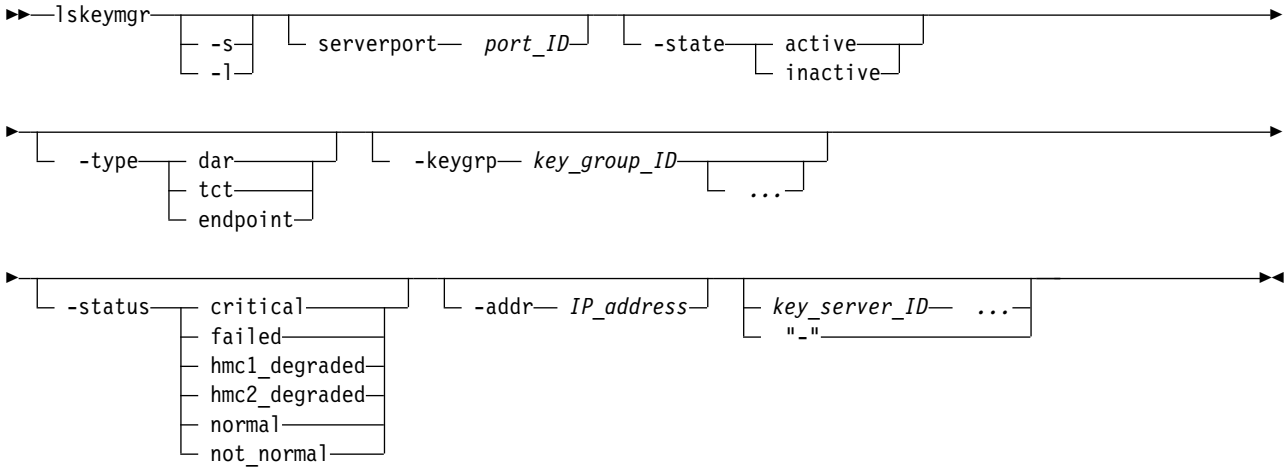
Endpoint
Encryption for IBM Fibre Channel Endpoint Security, which establishes authenticated communication and encryption of data in flight for Fibre Channel connections between an IBM Z host and the storage system. The connections are secured by Fibre Channel security protocols and key server authentication that uses communication certificates. If both the host and storage system use a connection with Fibre Channel ports that support encryption, the connection will transmit encrypted data between the ports.

name
Indicates the user-specified name that is used to identify the key group.

lskeymgr

The **lskeymgr** command displays a list of key servers that are on the storage complex and provides status information for those key servers.

The terms *key server* and *key manager* are closely related and sometimes used interchangeably. The key manager is the program that runs on a physical server and provides key services to key clients, such as the DS8000 system. Key services include generating, storing, and retrieving encryption keys. This command uses the term *key server* to indicate both the physical server and the key manager.



Parameters

-s (Optional) Displays the key server IDs.

-l (Optional) Displays the default output.

-serverport *port_ID*
(Optional) Displays the key servers that use the port ID that you specify. The key server port ID is 4 or 5 decimal characters. For example, 8100 is a valid port ID.

-state active | inactive
(Optional) Displays the key servers that are in the state that you specify.

-type dar | tct | endpoint
(Optional) Displays the registered key servers that use one of the following specified types of encryption.

dar Encryption for data at rest, which encrypts data that is stored on your storage system. This value is the default.

tct Encryption for transparent cloud tiering, which ensures that data is encrypted before it is transferred to cloud storage. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

endpoint
Encryption for IBM Fibre Channel Endpoint Security, which establishes authenticated communication and encryption of data in flight for Fibre Channel connections between an IBM Z host and the storage system. The connections are secured by Fibre Channel security protocols and key server authentication that uses communication certificates. If both the host and storage system use a connection with Fibre Channel ports that support encryption, the connection will transmit encrypted data between the ports.

-keygrp *key_group_ID | ...*
(Optional) Displays one or more key group IDs that are associated with the specified key manager. The ellipsis (...) indicates that, optionally, you can specify multiple values.

-status critical | failed | hmc1_degraded | hmc2_degraded | normal | not_normal
(Optional) Displays the status of the key server path. Only key servers that display the status of the specified key server path are displayed. The value `not_normal` displays for all key servers whose status is not "normal" or whose state is inactive.

-addr *IP_address*
(Optional) Displays the key server that uses the IP address you specify. The IP address can be an IPv4 address, an IPv6 address, or a DNS name.

key_server_ID ... | -
(Optional) Displays the key servers that use the ID or IDs you specify. To include multiple IDs, separate each ID with a blank space. For example, 1 2 3 4. The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the (-) dash option, this value can be read from standard input.

Example

Invoking the lskeymgr command to display the key server protocol.

```
dscli> lskeymgr -l
```

The resulting output

```
ID  State     Status  Addr                  Server  Key
                                            Port    Protocol
============================================================
1   active    normal  2001:0db8:85a3:08d3:  3801    KMIP
                      1319:8a2e:0370:7334
2   active    normal  2001:0db8::1428:57ab  3801    KMIP
3   inactive  -       ::ffff:9.11.56.78     3801    KMIP
```

Report field definitions

ID*
Indicates the key server identification number. For example, 4.

State
Indicates one of the following states of the key server displays:

active The key server is configured for a key exchange with the specified HMC.

inactive
The key server is configured, but does not exchange any key with the specified HMC.

Status
Indicates one of the following statuses the key server path:

critical
Only one Hardware Management Console (HMC) has access to the specified key server and it represents a potential single point of failure. This status `critical` replaces `hmc1_degraded` and `hmc2_degraded` on newer DS8000 systems.

Use the **showkeymgr** command with the –access parameter to determine the status of each HMC.

Note: A DS8000 system with only one HMC configured displays the status as `normal`.

degraded
Two or more HMCs have access to the specified key server, but at least one other HMC does not.

failed Neither HMC1 or HMC2 have access to the key server.

hmc1_degraded
HMC2 has access to the specified key server, but HMC1 does not.

Newer DS8000 systems display this status as `critical`.

Use the **showkeymgr** command with the –access parameter to determine the status of each HMC.

hmc2_degraded
HMC1 has access to the specified key server, but HMC2 does not.

Newer DS8000 models display this status as `critical`.

Use the **showkeymgr** command with the –access parameter to determine the status of each HMC.

normal
All HMCs have access to the specified key server.

"-" The dash (-) indicates that the specified key server is not an active key server.

Addr
Indicates the IP address of the key server.

Server Port
Indicates the key server port number, which is 4 or 5 decimal characters from 1 – 65535. For example, 8100.

Keyprotocol
The key server protocol that the storage system communicates with to provide key server management operations.

IPP The storage system communicates with the specified key server using the IBM Proprietary Protocol (IPP).

KMIP The storage system communicates with the specified key server using the Key Management Interoperability Protocol (KMIP).

Example

Invoking the lskeymgr command to display the type associated with the key server.

```
dscli> lskeymgr
```

The resulting output

```
| ID     state    reckeystate  reckeydate  datakeydate keyprotocol   type     name
| ==========================================================================================
| 1    accessible configured   02/21/2018   3/07/2018    KMIP      Endpoint endpoint_encr
```

Report field definitions

ID Indicates the key group ID.

state
Indicates one of the following states of the key group:

accessible
The key group is accessible if it is configured and the storage image has the key group from the key server for the specified key group.

inaccessible
The key group is inaccessible if the storage image was unable to obtain the key group from the key server.

unconfigured
The key group is unconfigured if it has not been configured.

rekeying
The key group is accessible and rekeying if it is configured and the storage image has the encryption key from the key server for the key group and is in the middle of rekeying.

reckeystate
Indicates one of the following states of the recovery key:

configured
A new recovery key was requested, verified, and authorized.

unconfigured
A recovery key was not created.

newkeyveripend
A new recovery key was requested but not verified.

newkeyauthpend
A new recovery key was requested and verified, but not authorized.

rekeyveripend
A new recovery key action was requested but not verified.

rekeyauthpend
A new recovery key action was requested and verified, but not authorized.

recovauthpend
A recover action was requested, but not authorized.

deconfauthpend
A deconfigure action was requested, but not authorized.

disabled
A recovery key was disabled, and the key group is used without a recovery key.

enableauthpend
An enable action was requested, but not authorized.

disableauthpend
A disable action was requested, but not authorized.

reckeydate
Indicates the date of the last recovery key creation.

datakeydate
Indicates the date of the last data key creation. If the key group is unconfigured, then any displayed date is to be considered erroneous data.

keyprotocol
Indicates the key server protocol that the storage system communicates with to provide key server management operations.

IPP
The storage system communicates with the specified key server using the IBM Proprietary Protocol (IPP).

KMIP
The storage system communicates with the specified key server using the Key Management Interoperability Protocol (KMIP).

type
Indicates the type of encryption that is used by the key group. The descriptions are described

DAR Encryption for data at rest, which encrypts data that is stored on your storage system. This value is the default.

TCT Encryption for transparent cloud tiering, which ensures that data is encrypted before it is transferred to cloud storage. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

Endpoint
Encryption for IBM Fibre Channel Endpoint Security, which establishes authenticated communication and encryption of data in flight for Fibre Channel connections between an IBM Z host and the storage system. The connections are secured by Fibre Channel security protocols and key server authentication that uses communication certificates. If both the host and storage system use a connection with Fibre Channel ports that support encryption, the connection will transmit encrypted data between the ports.

name
Indicates the user-specified name that is used to identify the key group.

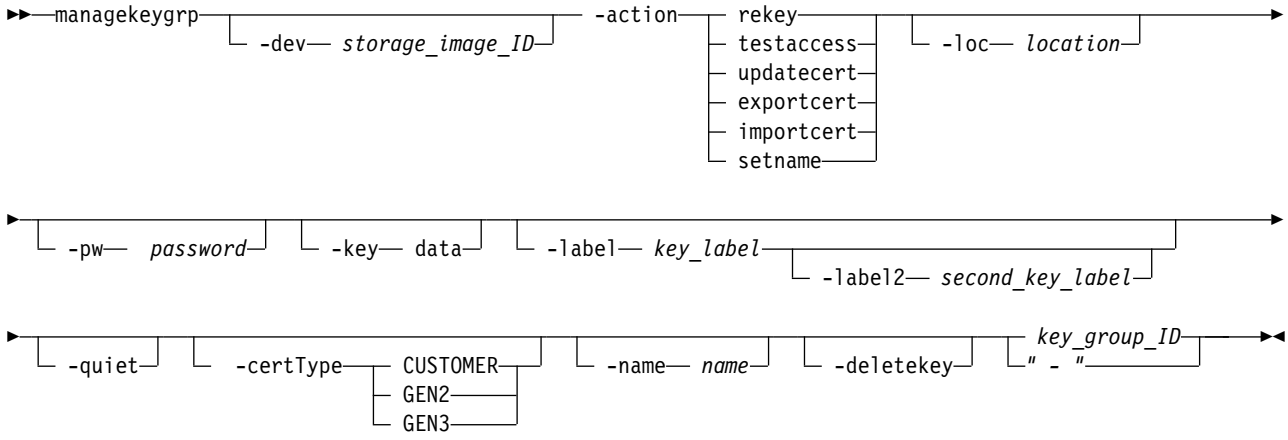
managekeygrp

The **managekeygrp** command manages a specified key group.

The encryption key servers use key labels to locate a specific encryption key that is stored in the key server. The dual platform key server supports a second label for the second key server platform. If the action is `rekey` or `updatecert` and no labels are specified, the existing labels are reused. If only one label is specified, that label is used for both key server platforms, replacing any existing labels. If two labels are specified, one label is used with the first key server platform, and the second label is used with the second platform. This process replaces any existing labels.

Note: The **managekeygrp** command cannot be run from the Embedded DS CLI window.

The connection to the encryption key that is stored on each key server is tested once every 8 hours. To test the key access for a specific key group on demand, use this command with the **–action testaccess** parameter. To view the results, use the **showkeygrp** command with the **–access** parameter.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. For example, IBM.2107-75FA120. The storage image ID is required if you do not specify a fully qualified key group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the storage system is aware of more than one storage image.

Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-action rekey | testaccess | updatecert | exportcert | importcert | setname
(Required) Specifies an action for the key group. The valid values are as follows.

Note: Starting from Release 8.1, DS8880 ships with GEN2 certificate enabled.

rekey
Specifies that a new key, as specified by the –**key** parameter, is generated.

testaccess
Specifies that access to the specified key group must be tested. Use the **showkeygrp –access** command to view the results.

updatecert
Specifies that the new key, which is specified by the –**key** parameter is generated. When the parameter processes, the key client certificate is updated to a NIST SP 800-131a compliant GEN2 certificate or GEN3 certificate that is used to authenticate with the key server.

Note: Before you use the `updatecert`, ensure that all key servers that are used by the DS8000 system contain the GEN2 trust anchor certificate.

exportcert
Specifies the option for exporting a DS8000 public certificate to a file on the DS CLI client. This option, available to users with administration authority, requires the **-loc** parameter.

Users with administration authority can export a DS8000 public certificate to a file on the DS CLI client.

importcert
Specifies the customer-defined certificate to import in Public-Key Cryptography Standards (PKCS) #12 format. The file name extension for PKCS #12 files is *.p12* or *.pfx*. This option does not update the current certificate. To update the current certificate, use the `updatecert` option.

When you use this option, specify the local file name of the customer-defined certificate with the **-loc** parameter and the password of the certificate with the **–pw** parameter.

setname
Specifies the option to specify a new name for the key group.

-loc *location*
(Optional) Specifies the location of the certificate file. For example, c:\mystore\trust.pem

The **-loc** parameter is required for the `exportcert` or `importcert` options.

-pw *password*
(Optional) Specifies the password of the certificate. For example, myPassword.

The **-pw** parameter is required for the `importcert` option.

The maximum length of customer-defined machine password is 128 characters. An example of a file name for a customer-defined filename is either 2107-75YZ123.mug or 2107-75YZ123.p12 and password is XyZABcPQRstu.

-key data
(Optional) Specifies that the key group data key is rewrapped with the `rekey` action. The key can be rewrapped with either new or existing key labels.

The **-key** parameter is required for the `rekey` or `updatecert` actions. It is optional for the `testaccess` action. This parameter specifies which key is the object of the specified action.

The data parameter specifies a random 256-bit encryption key that is generated by the key server and is used by the DS8000 system to wrap the key group. This key is identified by one of the key labels that is specified by the –**label** or –**label2** parameters.

-label *key_label*
(Optional) Specifies the label for the key group data key. This parameter is required when you specify the **-label2** parameter. You can enter a maximum of 64 ASCII characters for the label. The label is used by a key server to identify a specific data key that the DS8000 system uses to wrap the group key.

-label2 *second_key_label*
(Optional) Specifies the parameter that is used by a second key server to identify a specific data key that the DS8000 system uses to wrap the group key. You can enter a maximum of 64 ASCII characters for the label.

-quiet
(Optional) This parameter turns off the confirmation prompt for the command. This parameter is processed for `testaccess` only.

-certType CUSTOMER | GEN2 | GEN3
(Optional) Specifies the type of certificate. This parameter is optional if you specify either the `updatecert` or `exportcert` options.

CUSTOMER
Specifies the installed customer-defined certificate. If you specify the `updatecert` option, an IBM supplied GEN2 certificate or IBM supplied GEN3 certificate is updated to a customer-defined certificate. If you do not specify this option, the certificate defaults to an IBM GEN2 certificate. The parameter is optional for the `updatecert` and `exportcert` options.

The `exportcert` option is not supported by the **CUSTOMER** option.

Note: A customer-defined encryption certificate must meet the following requirements:

- The certificate must have an RSA public/private key pair with a single private key.
- The hashing algorithm must be SHA-256 or SHA-512.
- The public exponent must be 65537 or less.

- The public key size must be 2048 or 4096.
- The "Not Before" date must be before the current date.
- The "Not After" date must be after the current date.
- The key server ID must be in the valid range 1 - 4.
- The CN field value must be the storage facility image (SFI) ID for the storage system. For example, CN=2107-75NR699.
- For KMIP key groups that use a certificate user ID for Safenet key server local authentication, the UID field value must use the following format: DS8K-SFI_ID. For example, UID=DS8K-2107-75NR699.

GEN2
Specifies the SHA-256 digital signature encryption certificate that is included with the storage system. If you specify the `updatecert` option, the GEN2 certificate is applied. If you do not specify this option, the default is the IBM GEN2 certificate.

Users with administrative authority can use the `exportcert` option to export an IBM GEN2 certificate. If you do not specify this option, the certificate that is in use is exported.

GEN3
Specifies the SHA-512 digital signature encryption certificate, which includes the WWNN of the storage system and unique certificate serial number. GEN3 includes support for IBM Fibre Channel Endpoint Security, a feature that enables an IBM Z system to encrypt data and securely authenticate to the IBM DS8000 storage system. Before enabling IBM Fibre Channel Endpoint Security, update to GEN3 or user-defined certificates first.

Users with administrative authority can use the `exportcert` option to export an IBM GEN2 or IBM GEN3 certificate that is not enabled for IBM Fibre Channel Endpoint Security. If you do not specify this option, the certificate that is in use is exported.

-name
(Optional) Specifies the name that you specify to identify the key group.

-deletekey
(Optional) Specifies the option to use during the key group rekey action and updatecert action for the specified key group.

If you do not specify **-deletekey**, the key group rekey action and updatecert action are set (by default) to deactivate.

key_group_ID | -
(Required) Specifies the ID for the key group. The key group ID is a decimal number that ranges from 1 to N, where N is the maximum number of key groups that are supported by the storage system. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

```
dscli> managekeygrp –action updatecert -key data -certType GEN3 1
```

The resulting output

```
CMUC00472I managekeygrp: The certificate for key group 1 has been updated.
```

Example

```
dscli> managekeygrp –action importcert -loc /home/cert/75NH481/2107-75NH491_v3.mfg -pw
nPx6ZXXbZc2N 1
```

The resulting output

```
CMUC00489I managekeygrp: The certificate for key group 1 has been imported.
```

Example

```
dscli> managekeygrp –dev IBM.2107-75FA120 -action importcert -loc c:\hmcdir\2107-75ABCDE.mfg -pw
myPassword 1
```

The resulting output

```
managekeygrp: The certificate file has been imported successfully.
```

Example

```
dscli> managekeygrp –dev IBM.2107-75FA120 -action rekey -key data -label companyABC -label2
companyXYZ 1
```

The resulting output

```
CMUCFFFFFI managekeygrp: The rekey action has been performed for
key server key group 1.
```

managekeymgr

The **managekeymgr** command allows you to manage an existing encryption key server.

The DS8000 tests the access from each hardware management console (HMC) to each key server once every five minutes. To test access for a specific key manager on demand, use this command with the **–action testaccess** parameter. To view the results, use the **showkeymgr** with the **–access** parameter.

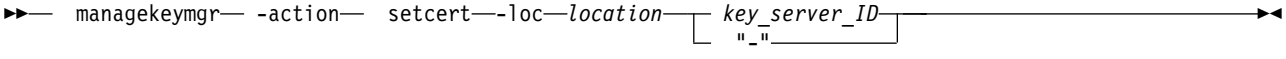
Note: The **managekeymgr** command cannot be run from the Embedded DS CLI window.

The terms *key server* and *key manager* are closely related and sometimes used interchangeably. The key manager is the program that runs on a physical server and provides key services to key clients, such as the DS8000 system. Key services include generating, storing, and retrieving encryption keys. This command uses the term *key server* to indicate both the physical server and the key manager.

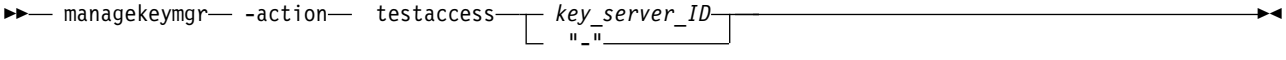


Each of the following sections shows the options for one of the listed **-action** parameters:

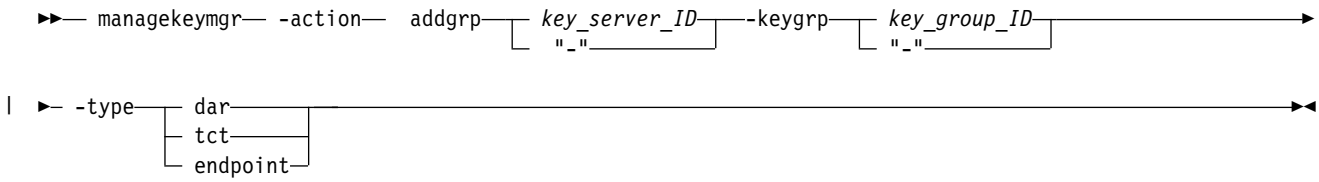
-action setcert



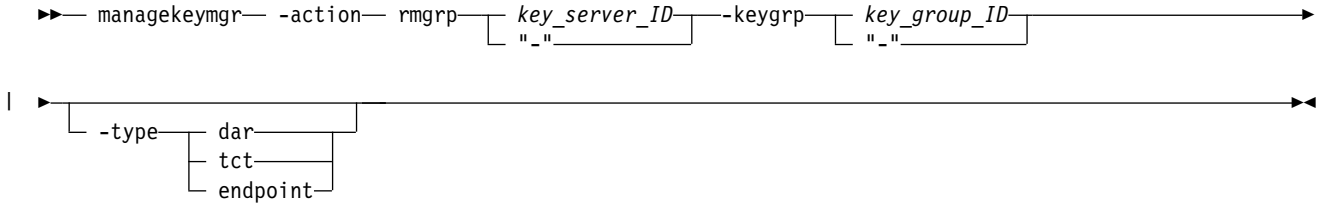
-action testaccess



-action addgrp



-action rmgrp



Action parameters for setcert

-action setcert
Specifies the replacement of the existing anchor certificate that is used in secure communications with the key server.

-loc *location*
(Required) Specifies the location of the certificate file to use as the trust anchor to authenticate communications to the specified key server. The certificate is in PEM or DER format. For example, `c:\mystore\trust.pem`

key_server_ID | **-**
(Required) Specifies the key server ID, in the range from 1 - *n*, where *n* is the number of key servers. Use the **showsp** command to determine the number of supported key servers.

Action parameters for testaccess

-action testaccess
Specifies to test (by pinging) the connection from each HMC to the specified key server. The results are displayed in the **showkeymgr** command by using the –access parameter.

key_server_ID | **-**
(Required) Specifies the key server ID, in the range from 1 - *n*, where *n* is the number of key servers. Use the **showsp** command to determine the number of supported key servers.

Action parameters for addgrp

-action addgrp
Specifies to add a key group to the specified key server.

key_server_ID | **-**
(Required) Specifies the key server ID, in the range from 1 - *n*, where *n* is the number of key servers. Use the **showsp** command to determine the number of supported key servers.

-keygrp *key_group_ID* | **-**
(Required) Specifies the ID for the key group. The key group ID is a decimal number that ranges from 1 to *n*, where *n* is the maximum number of key groups that are supported by the DS8000. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

-type dar | **tct** | **endpoint**
(Required) Specifies the type of encryption that is used by the key group.

dar Encryption for data at rest, which encrypts data that is stored on your storage system. This value is the default.

tct Encryption for transparent cloud tiering, which ensures that data is encrypted before it is transferred to cloud storage. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

endpoint
Encryption for IBM Fibre Channel Endpoint Security, which establishes authenticated communication and encryption of data in flight for Fibre Channel connections between an IBM Z host and the storage system. The connections are secured by Fibre Channel security protocols and key server authentication that uses communication certificates. If both the host and storage system use a connection with Fibre Channel ports that support encryption, the connection will transmit encrypted data between the ports.

Action parameters for rmgrp

-action rmgrp
Specifies to remove a key group to the specified key server.

key_server_ID | **-**
(Required) Specifies the key server ID, in the range from 1 - *n*, where *n* is the number of key servers. Use the **showsp** command to determine the number of supported key servers.

-keygrp *key_group_ID* | **-**
(Required) Specifies the ID for the key group. The key group ID is a decimal number that ranges from 1 to *n*, where *n* is the maximum number of key groups that are supported by the DS8000. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

-type dar | tct | endpoint
(Optional) Specifies the type of encryption that is used by the key group.

dar Encryption for data at rest, which encrypts data that is stored on your storage system. This value is the default.

tct Encryption for transparent cloud tiering, which ensures that data is encrypted before it is transferred to cloud storage. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

endpoint
Encryption for IBM Fibre Channel Endpoint Security, which establishes authenticated communication and encryption of data in flight for Fibre Channel connections between an IBM Z host and the storage system. The connections are secured by Fibre Channel security protocols and key server authentication that uses communication certificates. If both the host and storage system use a connection with Fibre Channel ports that support encryption, the connection will transmit encrypted data between the ports.

Example

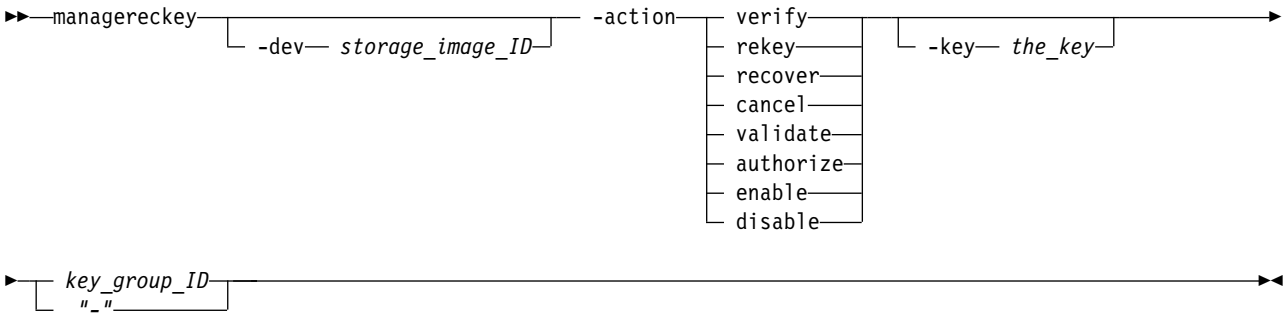
```
dscli> managekeymgr -action setcert -loc c:\mystore\trust.pem 4
```

The resulting output

```
The key server entry 4 successfully modified.
```

managereckey

The **managereckey** command is used to manage an existing encryption recovery key.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. For example, IBM.2107-75FA120. The storage image ID is required if you do not specify a fully qualified key group ID. It is also required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-action *verify* | *rekey* | *recover* | *cancel* | *validate* | *authorize* | *enable* | *disable*
(Required) Specifies the action to complete on the recovery key.

Verify
The *verify* option is for users with security administrator authority, and is the second step in creating an Encryption Recovery Key or re-creating an existing Encryption Recovery Key. Verify that you received the new Recovery Key from the first step (mkreckey, or managereckey with the -rekey option), by specifying that new key with the **-key** parameter. The next step requires the storage administrator to authorize the pending operation.

Rekey
The *rekey* option is for users with security administrator authority, and is the first step to reconfigure an existing Encryption Recovery Key. The existing Encryption Recovery Key is not required to start the rekey operation. The next step requires the security administrator to verify the new recovery key.

Recover
The *recover* option is for users with security administrator authority, and is the first step to using the Encryption Recovery Key to recover access to the encrypted data on the DS8000. The user specifies the Encryption Recovery Key with the **-key** parameter. The next step requires the storage administrator to authorize the recover operation.

Cancel
The *cancel* option is for users with either security administrator or storage administrator authority to cancel any verification or authorization pending steps. The existing Encryption Recovery Key is not required and no further steps are required.

Validate
The *validate* option is for users with security administrator authority, and is used to ensure that the existing Encryption Recovery Key is identical to the recovery key that is in the user's possession. The user specifies the Encryption Recovery Key in their possession with the **-key** parameter, but no further steps are required.

Authorize
The *authorize* option is for users with storage administrator authority, and is the final step of most recovery key operations. Once authorized, the pending recovery key operation is completed and any resulting changes to the DS8000 are started. The existing Encryption Recovery Key is not required and no further steps are required.

Enable
The *enable* option is for users with security administrator authority, and is the first step to enable the Encryption Recovery Key for the key group. The Recovery Key needs to be enabled only if it was previously disabled. The existing Encryption Recovery Key is not required. The next step requires the storage administrator to authorize the pending operation.

Disable
The *disable* option is for users with security administrator authority, and is the first step to disable the Encryption Recovery Key for the key group. The existing Encryption Recovery Key is not required. The next step requires the storage administrator to authorize the pending operation. After that, the group will operate without an Encryption Recovery Key, and will not be recoverable with a Recovery Key.

-key *the_key*
(Optional) Specifies the encryption recovery key. The encryption recovery key is a string of 64 hexadecimal characters.

For example, 01F3-45A7-8D12-B586-0123-4C67-891E-3586-01A3-45E7-8D12-3586-0123-45C7-8912-3B86.

key_group_ID | **-**
(Required) Specifies the key group ID for the encryption recovery key that you want to manage. The key group ID is a decimal number that ranges from 1 to N, where N is the maximum number of key groups that are supported by the DS8000. Use the **showsi** command to determine this maximum number. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

An invocation example

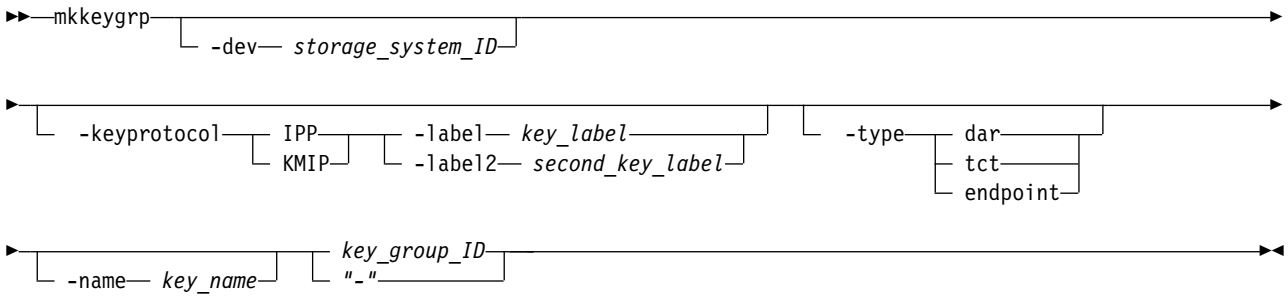
```
dscli> managereckey -dev IBM.2107-75FA120 -action verify
-key 0123-4567-8912-3586-0123-4567-8912-3586-0123-4567-8912-
3586-0123-4567-8912-3586 1
```

The resulting output

```
The Recovery Key for key group 1 has been verified,
authorization pending.
```

mkkeygrp

The **mkkeygrp** command creates an entry for the key group on the storage system.



Parameters

-dev *storage_system_ID*
(Optional) Specifies the storage system ID, which includes manufacturer, machine type, and serial number. For example, IBM.2107-75FA120.

Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-keyprotocol IPP | KMIP
(Optional) Specifies which key server protocols that the Security Key Lifecycle Manager supports to provide key server management operations for the storage system.

IPP Specifies that the storage system supports the IBM Proprietary Protocol (IPP) to communicate with the Security Key Lifecycle Manager for key server management operations. This value is the default.

KMIP Specifies that the storage system supports the Key Management Interoperability Protocol (KMIP) to communicate with the Security Key Lifecycle Manager for key server management operations.

-label *key_label*
(Required) Specifies the label for the key group data key. This parameter is required with the **-keyprotocol IPP** parameter. You can enter a maximum of 64 ASCII characters for the label.

-label2 *second_key_label*
(Optional) Specifies the second label for the key group data key. You can enter a maximum of 64 ASCII characters for the label.

-type dar | tct | endpoint
(Optional) Specifies the type of encryption that key group will use:

dar Encryption for data at rest, which encrypts data that is stored on your storage system. This value is the default.

tct Encryption for transparent cloud tiering, which ensures that data is encrypted before it is transferred to cloud storage. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

Note: If the **-type tct** option is specified, the data at rest encryption recovery key for the storage system must be disabled. After the key group is created, the recovery key can be enabled. For details about managing a recovery key, see the managereckey command.

endpoint
Encryption for IBM Fibre Channel Endpoint Security, which establishes authenticated communication and encryption of data in flight for Fibre Channel connections between an IBM Z host and the storage system. The connections are secured by Fibre Channel security protocols and key server authentication that uses communication certificates. If both the host and storage system use a connection with Fibre Channel ports that support encryption, the connection will transmit encrypted data between the ports.

-name *key_name*
(Optional) Specifies the user-specified name that is used to identify the key group.

key_group_ID | -
(Required) Specifies the ID for the key group. For key groups that use data at rest encryption, specify the numeral 1 as the ID. For key groups that use TCT encryption, specify a decimal number in the range 2 - 16. The default ID value is 1, which indicates a key group that is configured for data at rest encryption. Use the **showsi** command to determine this maximum number. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Example

```
dscli> mkkeygrp -dev IBM.2107-75FA120 -label MyCompany -label2 MyCompany2 1
```

The resulting output

```
The key server key group 1 has been created.
```

mkkeymgr

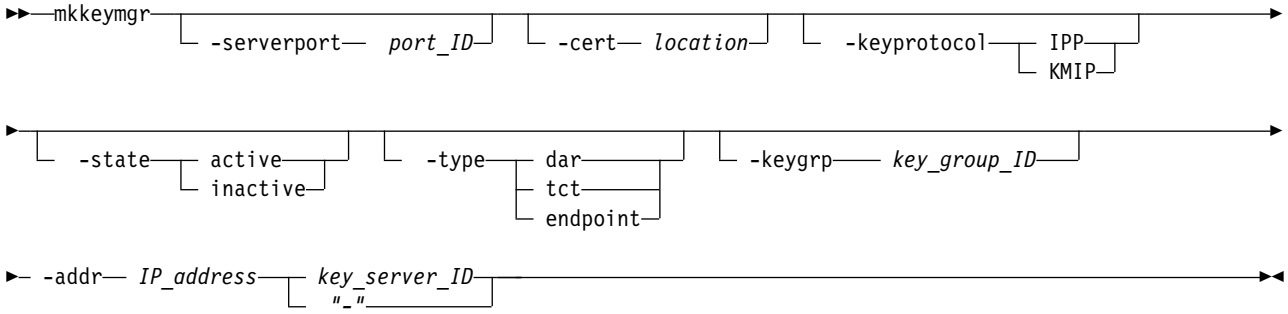
The **mkkeymgr** command creates an entry in the storage complex to access one or more key servers.

The DS8000 includes support for the industry-standard Key Management Interoperability Protocol (KMIP) for disk encryption. This support means that environments without IBM key management infrastructure can incorporate the DS8000 encryption feature into their environments that by using KMIP.

Note: The **mkkeymgr** command cannot be run from the Embedded DS CLI window.

The terms *key server* and *key manager* are closely related and sometimes used interchangeably. The key manager is the program that runs on a physical server and provides key services to key clients, such as the DS8000 system. Key services include generating, storing, and retrieving encryption keys. This command uses the term *key server* to indicate both the physical server and the key manager.

To use a key group, the key client on the DS8000 HMC must have access to two or more key servers. A given key server is identified by an IP address and an IP port number, where the IP port also determines whether the key server requires a security protocol (TLS or not TLS) in communications. The DS8000 key client uses a TLS protocol if, and only if, a certificate is provided. However, if a certificate is provided and the specified IP port is a non-TLS port, or if no certificate is provided and the specified IP port is a TLS port, then the DS8000 key client cannot connect to the key server.



Parameters

-serverport *port_ID*
(Optional) Specifies the key server port ID. The key server port ID is 4 or 5 decimal characters. For example, 8100 is a valid server port ID.

-cert *location*
(Optional) Specifies the location of the certificate file to use as a trust anchor to authenticate the certificate of the specified key server when using a TLS security protocol. If not specified, then only non-TLS protocols that do not require a trust anchor certificate are allowed. The certificate is in the PEM or DER format. For example, `C:\mystore\trust.pem`.

-keyprotocol IPP | KMIP
Specifies the key server protocol that the storage system communicates with to provide key server management operations.

IPP Specifies that the storage system communicates with the specified key server using the IBM Proprietary Protocol (IPP). This value is the default.

KMIP Indicates that the storage system communicates with the specified key server using the Key Management Interoperability Protocol (KMIP).

-state active | inactive
(Optional) Specifies the state of the key manager, where

active Specifies that this key manager is to be used to store encryption keys and that it should be checked periodically to verify the status of the key manager.

inactive
Specifies that this key manager is not to be used to store encryption keys and that it should not be checked periodically.

-type dar | tct | endpoint
(Optional) Specifies the type of encryption that is used by the key groups that are managed by the key manager.

dar Encryption for data at rest, which encrypts data that is stored on your storage system. This value is the default.

tct Encryption for transparent cloud tiering, which ensures that data is encrypted before it is transferred to cloud storage. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

endpoint
Encryption for IBM Fibre Channel Endpoint Security, which establishes authenticated communication and encryption of data in flight for Fibre Channel connections between an IBM Z host and the storage system. The connections are secured by Fibre Channel security protocols and key server authentication that uses communication certificates. If both the host and storage system use a connection with Fibre Channel ports that support encryption, the connection will transmit encrypted data between the ports.

-keygrp *key_group_ID* | **-**
(Optional) Specifies the ID for the key group. The key group ID is a decimal number that ranges from 1 to N, where N is the maximum number of key groups that are supported by the DS8000. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode. The default ID value is 1, which indicates a key group that is configured for data at rest encryption.

-addr *IP_address*
(Required) Specifies the IP address for the key manager. The IP address can be an IPv4 address, an IPv6 address, or a DNS name.

key_server_ID | **-**
(Required) Specifies the key server ID. The key server ID is a decimal number that ranges from 1 to N, where N is the maximum number of key servers that the DS8000 system can support. For example, 4.

Use the **showsp** command to determine this maximum number. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Example

```
dscli> mkkeymgr -cert /home/hscroot/keyppard.pem -keyprotocol KMIP -addr
keyppard.tuc1.stglabs.ibm.com 1
```

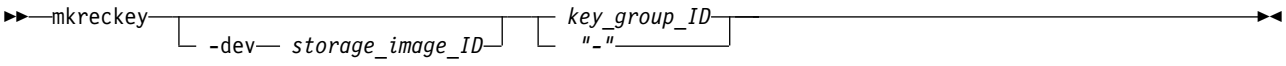
The resulting output

```
The key server 1 created successfully.
```

mkreckey

The **mkreckey** command is for users with security administrator authority. It is the first step in creating a new recovery key when no key currently exists.

If a key does exist, the security administrator must use the **managereckey -action** *rekey* command to rekey an existing key. The command returns the new key which the security administrator should copy to a safe place. The next step requires a security administrator to verify the new recovery key with the **managereckey** command.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. For example, IBM.2107-75FA120. The storage image ID is required if you do not specify a fully qualified key group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter will temporarily override any defined value for *devid* for the current command.

key_group_ID | **-**
(Required) Specifies the key group ID for the new recovery key that you want to create. The key group ID is a decimal number that ranges from 1 to N, where N is the maximum number of key groups supported by the DS8000. Use the **showsi** command to determine this maximum number. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

An invocation example

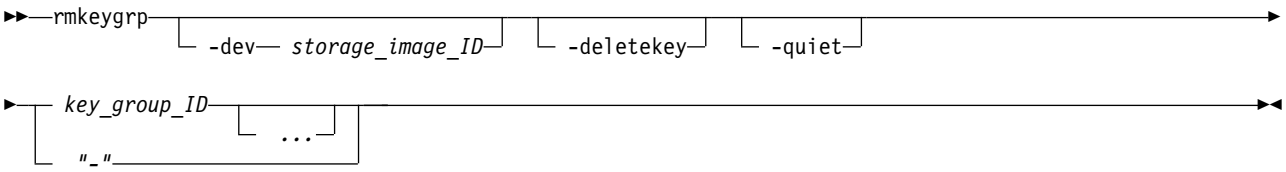
```
dscli> mkreckey -dev IBM.2107-75FA120 1
```

The resulting output

```
The Recovery Key 0123-4567-8912-3586-0123-4567-8912-3586-01
23-4567-8912-3586-0123-4567-8912-3586 for encryption group 1
has been created, verification pending.
```

rmkeygrp

The **rmkeygrp** command removes an entry for the key server group on a specified storage image.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified key group ID. The ID is also required if you do not set the *devid* variable in your profile or through the **setenv** command, and the storage system is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-deletekey
(Optional) Specifies the option to use during the key group delete action from the specified key server.

If you do not specify the **-deletekey** option, the key group delete action is set (by default) to deactivate.

-quiet
(Optional) Turns off the removal confirmation prompt for this command.

key_group_ID
(Required) Specifies the ID for the key group. The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

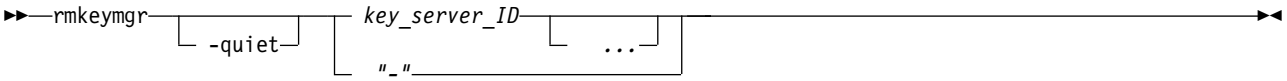
```
dscli> rmkeygrp -dev IBM.2107-75FA120 1
```

The resulting output

```
Are you sure you want to delete the key server group 1? y/n y
The key server group 1 has been deleted.
```

rmkeymgr

The **rmkeymgr** command removes a key server entry on the storage complex.



Parameters

-quiet
(Optional) Turns off the removal confirmation prompt for this command.

key_server_ID ... | -
(Required) Deletes the key servers with the specified key server ID. You must separate multiple IDs or ID ranges with a space between each value. The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

An invocation example

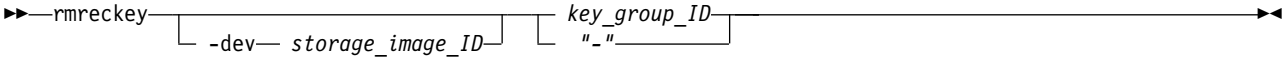
```
dscli> rmkeymgr -quiet 1 2
```

The resulting output

```
The key server 1 has been deleted.
The key server 2 has been deleted.
```

rmreckey

The **rmreckey** command allows you to remove an encryption recovery key.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial

number. For example, IBM.2107-75FA120. The storage image ID is required if you do not specify a fully qualified key group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter will temporarily override any defined value for *devid* for the current command.

key_group_ID | **-**
(Required) Specifies the key group ID for the encryption recovery key that you want to deconfigure. The key group ID is a decimal number that ranges from 1 to N, where N is the maximum number of key groups supported by the DS8000. Use the **showsi** command to determine this maximum number. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

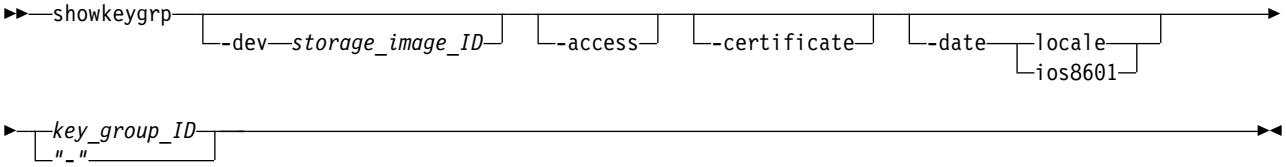
```
dscli> rmreckey -dev IBM.2107-75FA120 1
```

The resulting output

```
The Recovery Key for key group 1 will be de-configured after authorization.
```

showkeygrp

The **showkeygrp** command displays detailed information for a specified key group, including certificate information.



Parameters

-dev *storage_image_ID*
(Optional) Displays the properties for the specified storage image ID, which includes manufacturer, machine type, and serial number. For example, IBM.2107-75FA120. The storage image ID is required if you do not specify a fully qualified key group ID. Do not set the devid value in your profile or through the **setenv** command, and the HMC is aware of more than one storage image.

Using the **-dev** parameter temporarily overrides any defined value for devid for the current command.

-access
(Optional) Displays output with dates and times that the specified key group last assessed the encryption key on the key server.

-date locale | **iso8601**
(Optional) Specifies the date and time format as follows:

locale
Specifies the locale format as set with the **-local** parameter, or by default, from the OS locale setting.

iso8601
Specifies the ISO 8601 date/time format as follows:

```
yyyy-MM-dd'T'HH:mm:ssZ
```

where:

yyyy the year

MM the month (01-12)

dd the day (01-31)

HH the hour (00-23)

mm the minutes (00-59)

ss the seconds (00-59)

Z the time zone offset from UTC [-HHmm | +HHmm]

The default value is *locale*.

key_group_ID | **-**
(Required) Specifies the ID for the key group. The key group ID is a decimal number that ranges from 1 to *n*, where *n* is the maximum number of key groups supported by the DS8000. Use the **showsi** command to determine this maximum number. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the showkeygrp command to display key server information

```
dscli> showkeygrp -l
```

The resulting output

```
ID                              1
numranks                        1
numpools                        1
state                           accessible
reckeystate                     configured
reckeydate                      01/14/2016 15:05:30 MST
datakeydate                     01/14/2016 15:05:30 MST
grpstatus                       critical
mgrstatus                       normal
label                           CompanyABC
label2                          CompanyABC
certificate                     GEN2
certificate valid not before    02/03/2018 14:58:37 CST
certificate valid not after     01/27/2019 14:58:37 CST
certificate issuer              C=US, ST=AZ, L=Tucson, O=IBM, OU=Storage
certificate subject             C=US, ST=AZ, L=Tucson, O=IBM, OU=Storage
certificate install state       Imported, Ready to Update
| endpoint enabled              Yes
pending cert valid not before   05/08/2017 16:57:13 CST
pending cert valid not after    05/08/2018 16:57:13 CST
pending cert issuer             CN=2107-75NR641,O=Default Company Ltd,
                                    L=Default City,
                                ST=ShangHai,C=CN
pending cert subject            CN=2107-75NR641,O=Default Company Ltd,
                                    L=Default City,
                                ST=ShangHai,C=CN
| pending endpoint enabled      Yes
uuid                            128A829793D8A391ZZAF8X95554ACF813B58B0C98X2
                                B0DECF3067XXD4D161776
keyprotocol                     KMIP
| type                         ENDPOINT
| name                         ENDPOINT_1
```

Report field definitions

ID Indicates the key group ID.

numranks
Indicates the number of ranks that are configured in the specified key group.

numpools
Indicates the number of extent pools that are configured in the specified key group.

state Indicates one of the following states of the key group:

accessible
The key group is accessible if it is configured. The storage image obtained the encryption key from the key server for the key group.

inaccessible
The key group is inaccessible if the storage image was unable to obtain the encryption key from the key server.

unconfigured
The key group is unconfigured if it has not been configured.

rekeying
The key group is accessible and rekeying if it has been configured. The storage image obtained the encryption key from the key server for the key group and is in the process of rekeying.

reckeystate
Indicates one of the following states of the recovery key:

configured
A new recovery key was requested, verified, and authorized.

unconfigured
A recovery key has not been created.

newkeyveripend
A new recovery key has been requested but not verified.

newkeyauthpend
A new recovery key has been requested and verified, but not authorized.

rekeyveripend
A new recovery key action has been requested but not verified.

rekeyauthpend
A new recovery key action has been requested and verified, but not authorized.

recovauthpend
A recover action has been requested, but not authorized.

deconfauthpend
A deconfigure action has been requested, but not authorized.

disabled
A recovery key has been disabled, and the key group is used without a recovery key.

enableauthpend
An enable action has been requested, but not authorized.

disableauthpend
A disable action has been requested, but not authorized.

reckeydate
The date and time of the last recovery key creation.

datakeydate
The date and time of the last data key creation. If the key group is unconfigured, then any displayed date is to be considered erroneous data.

grpstatus
Indicates one of the following access statuses of the encryption key.

Note: The **grpstatus** parameter is refreshed by either a background process that runs every 8 hours or after the **managekeygrp** command with -action testaccess processes. The following values that display depend on either process to determine when the **grpstatus** parameter is refreshed.

For example, if all key servers were recovered from an abnormal status, the **grpstatus** parameter displays `failed` because the status reflects either process that last completed unless a background process has started or the **managekeygrp** command with -action testaccess has processed again.

In addition, when a key group is created, the status always displays a "-" to indicate that the background process has not started or the **managekeygrp** command with -action testaccess has not processed.

critical
Indicates that the key group has access to the encryption key on a single key server and it represents a potential single point of failure.

degraded
Indicates that the key group has access to the encryption key on two or more key servers, but not all key servers.

failed Indicates that the key group does not have access to the encryption key on any key server.

normal
Indicates that the key group has access to the encryption key on all key servers.

"-" Indicates by the dash (-) that the key group state is either unconfigured, rekeying, or was created but the background process has not started or the **managekeygrp** command with -action testaccess has not processed.

mgrstatus
Indicates one of the following values of the key server path access status:

critical
Indicates that at least one key server for this key group reported an access status of normal or critical.

Note: A DS8000 system with only one HMC configured displays status as `normal`.

degraded
Indicates that at least two key servers for this key group reported an access status of `normal` or `degraded.`

failed Indicates that all key servers for this key group reported an access status of `failed`.

normal
Indicates that all key servers for this key group reported an access status of `normal`.

"-" Indicates by the dash (-) that the key group state is either unconfigured or all key servers report a state of inactive.

label Indicates the label that is used by the key servers to locate a specific encryption key. Example MyCompany

label2 Indicates the second label that is used by the key servers to locate a specific encryption key.Example MyCompany2

certificate
Indicates the generation of the currently activated key client certificate that is used in secure communications with the key servers. The following values are possible:

CUSTOMER
Indicates the installed customer-defined certificate.

GEN1 Indicates that GEN1 is the legacy certificate that is not compliant with NIST SP 800-131a.

GEN2 Indicates that GEN2 is the NIST SP 800-131a-compliant certificate.

GEN3 Indicates that GEN3 is the NIST SP 800-131a-compliant certificate.

certificate valid not before
Indicates that the certificate is not valid before a certain date.

certificate valid not after
Indicates that the certificate is not valid after a certain date.

certificate issuer
Identifies the certificate issuer.

certificate subject
Identifies the certificate subject.

endpoint enabled
Identifies whether the IBM Fibre Channel Endpoint Security is enabled for the certificate.

certificate install state
Indicates one of several installation states:

- Not Imported
- Imported, Ready to Update
- Updated, In Use

pending cert valid not before
Indicates that the pending certificate is not valid before a certain date.

pending cert valid not after
Indicates that the pending certificate is not valid after a certain date.

pending cert issuer
Identifies the issuer of the pending certificate.

pending cert subject
Identifies the subject of the pending certificate.

pending endpoint enabled
Indicates whether IBM Fibre Channel Endpoint Security is enabled for the pending certificate.

UUID Indicates the Universally Unique Identifier (UUID). The KMIP key server generates an encryption key and assigns an UUID to it. The storage system then uses the UUID to identify the encryption key on the KMIP key server.

keyprotocol
Indicates one of the following values for the key protocol server:

IPP Indicates that the storage system communicates with the specified key server using the IBM Proprietary Protocol (IPP).

KMIP Indicates that the storage system communicates with the specified key server using the Key Management Interoperability Protocol (KMIP).

type The type of encryption that is used by the key group:

DAR Encryption for data at rest, which encrypts data that is stored on your storage system. This value is the default.

TCT Encryption for transparent cloud tiering, which ensures that data is encrypted before it is transferred to cloud storage. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

Endpoint
Encryption for IBM Fibre Channel Endpoint Security, which establishes authenticated communication and encryption of data in flight for Fibre Channel connections between an IBM Z host and the storage system. The connections are secured by Fibre Channel security protocols and key server authentication that uses communication certificates. If both the host and storage system use a connection with Fibre Channel ports that support encryption, the connection will transmit encrypted data between the ports.

Name Indicates the user-specified name used to identify the key group.

Example

Invoking the showkeygrp command to display certificate information

```
dscli> showkeygrp -certificate 1
```

The resulting output

```
ID                      1
certificate             CUSTOMER
version                 3
serial number           00:cf:86:54:6e:26:c7:c2:01
issuer                  O=ibmDisk,C=US
not valid before        08/24/2018 15:00:02 MST
not valid after         08/11/2032 17:36:55 MST
subject                 CN=2107-75NR641,O=ibmDisk,C=US,uid=DS8K-2107-75NR641
WWNN                    OTHERNAME=50:05:07:63:09:SF:C5:EE
size                    2048
signature algorithm     SHA256withRSA
SHA-256 fingerprint     1b:9d:7e:71:62:c9:80:4c:30:d8:4c:7c:6c:b0:c5:ac:4d:6f:4b:8c:ef:a7:f8:
                        ec:5e:fa:b6:20:c8:cf:65:93
public key algorithm    RSA
public key exponent     65537
authority key id        37:8a:3d:1b:8b:0c:1e:f2:c2:c5:ed:ea:8c:65:a9:07:fd:c2:a9:e4
subject key id          ad:8f:e8:ee:bd:e1:48:4f:c5:14:68:7b:52:a7:0d:c2:41:87:d7:16
pem format                 -------BEGIN CERTIFICATE-------
MIIDqjCCApKgAwIBAgIJAM+DVD4mycIBMA0GCSqGSIb3DQEBCwUAMB8xCzBJBgNVBAYTAlVTMRAwDgYDVQ
QKEwdpYm1EaXNrMB4XDTE4MDgyNDIyMDAwMloXDTMyMDgxMjAwMzY1NVowWTEhMB8GCgmSJomT8ixkBQE
TEURTOEstMjEwNy03NU5SNjQxMQswCQYDVQQGEwJVUzEQMA4GA1UEChMHaWJtRGlzazEVMBMGA1UEAxMMM
jEwNy03NU5SNjQxMIIBIjANCgkqhkiG9w0BAQEFAAOCAQ7AMIIBCgKCAQEArbVmVAIIbaARrMeIobZhgQe
hw9VA16Van7a15faaEJOyWg02oYts7pwC86B7JyzEYAgN6Gb/slY1YgW0m18axJcgbhW856DK+K9A23NSr
YHgv+Dqb7Qp+R6TD7jO1wv/ecY7nYAPgc4h5YWWO6+iMbnV7oHZkFRLSVdqCT2ntThupD2UzQBNLoSax9bk
+kkeU7DH/zCL1SjoHF/pbhzg1yIkuVLUkZGVGY1z0RLVa/gMTNQAwNnIwJnivswL83M17Wudq76ug6xS4zPj
FCBbCOMS0Nf/MhH68UE/lv4r63rJ6IWXwf2iadBM+xzJbr4tjHHX9MW8NkdiVOzW5dX0lwIDAQABo4GuMIGrMEc
GA1UdIwRAMD6AFDeKPRuLDB7ywsXt6oxlqQf9wqnkoSOkITAfMQswCQYDVQQGEwJVUzEQMA4GA1UEChMHaWJt
RGlza4IBADAdBgNVHQ4EFgQUrY/o7r3hSE/FFGh7UqcNwkGH1xYwDgYDVR0PAQH/BAQDAgWgMDEGA1UdEQQqM
CigJgYJKoZIhv1iAQEBoBMFzUwOjA1OjA3OjYzOjA5OkZGOkM1OkREMA0GCSqGSIb3DQEBCwUAA4IBAQBiCLl
7U4j96rUIU85+udvY6u+jfj9kbvIViZHQ/gJYP46Z0Om5iMmQgfV2YBJPOa96RCtDS5gJj/7xU5P7jTxfcMoUWU
bge37oDHy7QSG3of587Izf2quZskZOPUhjylzFgIQw89YuXbA6ivsneKjuAXAv1IIs3SiD7RE3I3PYv9X9S5R35e
tmJN9WlF+DKhI2uT92S91xTfDtXNm0PvVUJbGKgOh915rCkezbpltO/CvGUiw6cwufnva3rNjPL84a7gy9REY0oe
sW5JW4KpwMhfx9gr4aGNj/caE70gZ4N7XtGzXpsTB3LqOEESCU1sg0VXBwQXpXZjHiJu11HpzI
--------END CERTIFICATE--------
```

Report field definitions

ID Indicates the ID of the certificate.

certificate
Indicates the type of certificate.

version
Indicates the version of the certificate.

serial number
Indicates the serial number that uniquely identifies the certificate.

issuer
Indicates the entity that verified the information and signed the certificate.

not valid before
Indicates the earliest time and date on which the certificate is valid.

not valid after
Indicates the time and date past which the certificate is no longer valid.

subject
Indicates the entity a certificate belongs to: a machine, an individual, or an organization.

WWNN
Indicates the world wide node name that is associated with the certificate.

size Indicates the key size of the certificate.

signature algorithm
Indicates the cryptographic algorithm used to sign the certificate.

SHA-256 fingerprint
Indicates the unique identifier of the certificate.

public key algorithm
Indicates the algorithm that is used to encrypt a pair of keys known as a public key and a private key (a public key pair), which are associated with an entity that needs to authenticate its identity electronically or to sign or encrypt data. Each public key is published and the corresponding private key is kept secret. Data that is encrypted with the public key can be decrypted only with the corresponding private key.

public key exponent
Indicates the value that is used to shorten the key size for encryption and signature verification.

authority key id
Indicates the ID that identifies the public key that is being used to verify the signature on this certificate.

subject key id
Indicates the ID that identifies the public key that is being certified.

pem format
Indicates the privacy enhanced mail (PEM) format for the certificate.

Example

Invoking the showkeygrp command to display encryption key access per key server (KeyMgr)

```
dscli> showkeygrp -access 1
```

The resulting output

```
ID   KeyMgr   LastAccess   LastSuccess              LastFailure
-------------------------------------------------------------------------
1    1        success      2013-12-11T06:37:46-700  2013-04-26T16:23:18-700
1    2        failure      2013-11-29T21:22:09-700  2013-12-11T06:37:47-700
```

Report field definitions

ID Identifies the key group ID in ranges from 1 to the number of key groups. Use the**showsi** command to determine the number of supported groups.

KeyMgr
Indicates the key server ID in ranges from 1 to the number of key servers. Use the **showsp** command to determine the number of supported key servers.

LastAccess
Indicates the result of the last access attempt. The following values are possible:

success
The LastSuccess date is more recent than the LastFailure date.

failure
The LastFailure date is more recent than the LastSuccess date.

inactive
The key server is in an inactive state. The encryption key access test to this key server cannot be performed.

A "-" (dash) displays if the specified encryption key server has never attempted to access the encryption key.

LastSuccess
Indicates the date and time of the last successful attempt of the specified key group to access the encryption key on the key servers. The output displays in ISO 8601 data/time format.

```
yyyy-MM-dd'T'HH:mm:ssZ
```

where:

yyyy The year

MM The month (01-12)

dd The day (01-31)

HH The hour (00-23)

mm The minutes (00-59)

ss The seconds (00-59)

Z The time zone offset from UTC [-HHmm | +HHmm]

"-" A "-" (dash) displays if the specified key server has never successfully accessed the encryption key.

LastFailure
Indicates the date and time of the last failed attempt of the specified key group to access the encryption key on the key servers. The output displays in ISO 8601 data/time format.

```
yyyy-MM-dd'T'HH:mm:ssZ
```

where:

yyyy The year

MM The month (01-12)

dd The day (01-31)

HH The hour (00-23)

mm The minutes (00-59)

ss The seconds (00-59)

Z The time zone offset from UTC [-HHmm | +HHmm]

"-" A "-" (dash) displays if the specified key server has never failed to access the encryption key.

showkeymgr

The **showkeymgr** command displays detailed properties of a specified key server entry. This command also displays when the hardware management console (HMC) last accessed the specified key server.

The terms *key server* and *key manager* are closely related and sometimes used interchangeably. The key manager is the program that runs on a physical server and provides key services to key clients, such as the DS8000 system. Key services include generating, storing, and retrieving encryption keys. This command uses the term *key server* to indicate both the physical server and the key manager.



Parameters

-access
(Optional) Displays a table that includes dates and times that each HMC last accessed the specified key server.

key_server_ID
(Required) Specifies the key server ID, in the range from 1 - *n*, where *n* is the number of key servers. Use the **showsp** command to determine the number of supported key servers.

Example

Invoking the showkeymgr command to display key server information

```
dscli > showkeymgr 1

Date/Time: January 3, 2018 11:37:07 AM MST IBM DSCLI Version: 0.0.0.0 DS:
```

The resulting output

```
ID          1
State       active
Status      normal
addr        keyhler.tuc.stglabs.ibm.com
port        5697
protocol    tls
keyprotocol KMIP
certificate keyhler.pem
type        DAR
keygrp      1
```

Report field definitions

ID Indicates the key server ID, in the range from 1 - *n*, where *n* is the number of key servers. For example, 1.

Use the **showsp** command to determine the number of supported key servers.

State Indicates the key activation state with the following values:

active
Indicates that the key server is configured for a key exchange with the specified HMC.

inactive
Indicates that the key server is configured, but will not exchange any key with the specified HMC.

Status Indicates the status of the key path with one of the following values:

critical
Indicates that only one HMC has access to the specified key server and it represents a potential single point of failure. This status replaces `hmc1_degraded` and `hmc2_degraded` on newer DS8000 systems.

Note: A DS8000 system with only one HMC configured displays status as `normal`.

degraded
Indicates that two or more HMCs have access to the specified key server, but at least one HMC does not.

failed
Indicates that neither HMC1 or HMC2 have access to the specified key server.

hmc1_degraded
Indicates that HMC2 has access to the specified key server, but HMC1 does not.

Newer DS8000 systems display this status as `critical`.

hmc2_degraded
Indicates that HMC1 has access to the specified key server, but HMC2 does not.

Newer DS8000 systems display this status as `critical`.

normal
Indicates that all HMCs have access to the specified key server.

"-"
Indicates by the dash (-) that the specified key server is not an active key server.

addr Indicates that the key server IP address is in an IPv4 or IPv6 format.

port Indicates the key server port number, which is 4 or 5 decimal characters from 1 – 65535. For example, 8100.

protocol
Indicates the connection protocol, where:

none
Indicates that no encryption protocol is used.

tls
Allows protocols SSLv3, TLSv1.0, TLSv1.1, and TLSv1.2.

Note: The actual TLS protocol that is used depends also on whether the DS8000 system was set to NIST 899131a-compliant.

keyprotocol
Indicates one of the following values for the key protocol server:

IPP Indicates that the storage system communicates with the specified key server using the IBM Proprietary Protocol (IPP).

KMIP Indicates that the storage system communicates with the specified key server using the Key Management Interoperability Protocol (KMIP).

certificate
Indicates the certificate file. A dash (-) is displayed if no certificate was specified.

type The type of encryption that is used by the key groups that are managed by the key manager.

DAR Encryption for data at rest, which encrypts data that is stored on your storage system. This value is the default.

TCT Encryption for transparent cloud tiering, which ensures that data is encrypted before it is

transferred to cloud storage. The data remains encrypted in cloud storage and is decrypted after it is transferred back to the storage system.

Endpoint
Specifies encryption for IBM Fibre Channel Endpoint Security, which provides authentication of participating hosts and storage devices to ensure a secure transmission of encrypted data using industry methods.

keygrp
Indicates the key group ID that is associated with this key manager.

Example

Invoking the showkeymgr command to display access information about the specified key server.

```
dscli > showkeymgr -access 1
```

The resulting output

```
Date/Time: December 11, 2013 8:57:44 AM MST IBM DSCLI Version: 7.7.20.524 DS:

ID  hmc    LastAccess  LastSuccess               LastFailure
-------------------------------------------------------------------------
1   hmc1   success     2017-12-11T06:37:46-700   2017-04-26T16:23:18-700
1   hmc2   failure     2017-11-29T21:22:09-700   2017-12-11T06:37:47-700
```

Report field definitions

ID Indicates the key server ID, in the range from 1 - *n*, where *n* is the number of key servers. For example, 1.

Use the **showsp** command to determine the number of supported key servers.

hmc Indicates the hmc on which access was attempted. The possible values are hmc1 (primary) or hmc2 (secondary).

LastAccess
Indicates the result of the last access attempt. The possible values include `success` and `failure`.

A "-" (dash) displays if the HMC has never attempted to access the specified key server.

LastSuccess
Indicates the date and time of the last successful attempt by the HMC to access the specified key server. The output displays in ISO 8601 data/time format.

```
yyyy-MM-dd'T'HH:mm:ssZ
```

where:

yyyy the year

MM the month (01-12)

dd the day (01-31)

HH the hour (00-23)

mm the minutes (00-59)

ss the seconds (00-59)

Z the time zone offset from UTC [-HHmm | +HHmm]

"-" A "-" (dash) displays if the HMC has never successfully accessed the specified key server.

LastFailure
Indicates the date and time of the last failed attempt by the HMC to access the specified key server. The output displays in ISO 8601 data/time format.

```
yyyy-MM-dd'T'HH:mm:ssZ
```

where:

yyyy the year

MM the month (01-12)

dd the day (01-31)

HH the hour (00-23)

mm the minutes (00-59)

ss the seconds (00-59)

Z the time zone offset from UTC [-HHmm | +HHmm]

"-" A "-" (dash) displays if the HMC has never failed to access the specified key server.

User access security commands

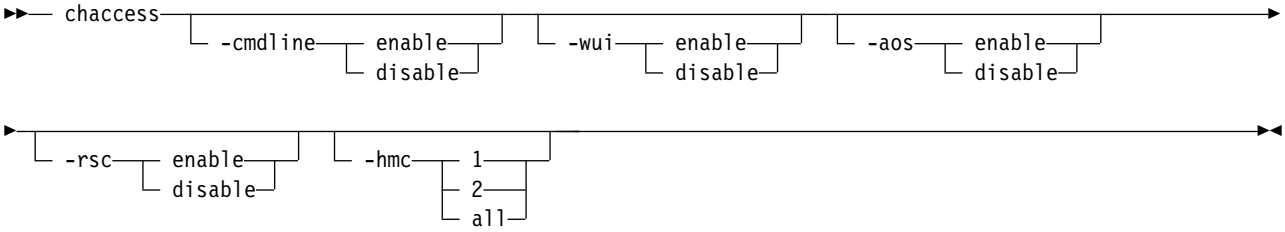
User access commands are used to manage network security on the DS8000 system by using the DS CLI.

The following user access security commands are available:

- **chaccess** - The **chaccess** command allows for changing one or more access settings of a Hardware Management Console.
- **lsaccess** - The **lsaccess** command displays the access settings of a Hardware Management Console.
- **manageaccess** - The **manageaccess** command manages the security protocol access settings of a Hardware Management Console (HMC) for all communications to and from the DS8000 system.
- **showaccess** - The **showaccess** command displays the access properties of a specified HMC.

chaccess

The **chaccess** command changes one or more access settings of a hardware management console (HMC). Only users with administrator authority can access this command.



Parameters

-cmdline enable | disable
(Optional) Specifies whether you want to enable or disable the command line shell access to the HMC through the Internet or a VPN connection. This command affects service access only and does not change access to the system through the DS Command Line Interface. At least one parameter **-cmdline**, **-wui**, or **-aos** must be specified.

-wui enable | disable
(Optional) Specifies whether you want to enable or disable the Web User Interface (WUI) access on

the HMC through the Internet or a VPN connection. This command affects service access only and does not change access to the storage system through the GUI. At least one parameter **-cmdline**, **-wui**, or **-aos** must be specified.

-aos enable | disable
(Optional) Specifies whether you want to enable or disable Assist On-Site (AOS) connection to the HMC for IBM service. At least one parameter **-cmdline**, **-wui**, or **-aos** must be specified.

-rsc enable | disable
(Optional) Specifies whether you want to enable or disable the remote support center (rsc) service for remote access. At least one parameter **-cmdline**, **-wui**, **-aos**, or **-rsc** must be specified.

-hmc 1 | 2 | all
(Optional) Specifies to which HMC you want the access settings to apply. hmc 1 specifies the primary HMC, and hmc 2 specifies the secondary HMC. The default value all specifies the primary HMC on a single HMC system, and both the primary and secondary HMCs on a dual HMC system.

Example

Invoking the chaccess command

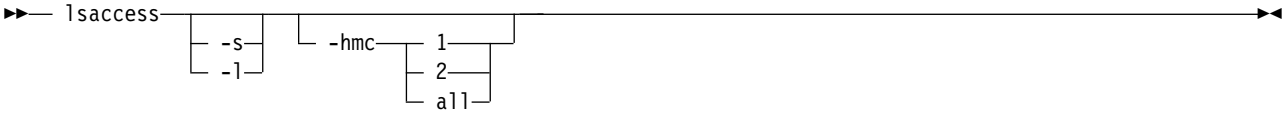
```
dscli> chaccess –cmdline enable –wui enable -hmc 1
```

The resulting output

```
hmc1 successfully modified.
```

lsaccess

The **lsaccess** command displays the access settings of the primary and backup hardware management consoles (HMCs).



Parameters

-s
(Optional) Displays only attributes that are identified as short output. You cannot use the **-s** and the **-l** parameters together.

-l
(Optional) Displays default output plus extra attributes that are identified as long output. You cannot use the **-l** and the **-s** parameters together.

-hmc 1 | 2 | all
(Optional)

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

Invoking the lsaccess command to display a list of access properties

```
dscli> lsaccess -l
```

The resulting output

```
hmc             hmc1
cmdline         enabled
wui             enabled
aos             disabled
vpn             disabled
rsc             enabled
```

Report field definitions

hmc*
Displays the HMC identifier.

cmdline
Indicates whether the command line interface is enabled or disabled for service access through an Internet or VPN connection. If the client is not connected to an HMC, a dash (-) indicates that the information is not available.

This field is for service access only and is not related to access to the storage system through the DS CLI.

wui
Indicates whether the Web user interface (WUI) is enabled or disabled for service access through an Internet or VPN connection. If the client is not connected to an HMC, a dash (-) indicates that the information is not available.

aos
Indicates whether the setting for Assist On-Site (aos) is enabled or disabled for connection to the HMC for IBM service. If the client is not connected to an HMC or if the query is not supported by the server, a dash (-) indicates that the information is not available.

vpn+
Indicates the connection status (enabled or disabled) of the outbound VPN connection. A status of enabled means that a VPN connection exists, and a status of disabled means that a VPN connection does not exist. Indicates whether the modem is enabled or disabled for access through an Internet or VPN connection. If the client is not connected to an HMC, a dash (-) indicates that the information is not available.

Note: Starting with DS8000 Release 8.0, VPN is no longer available as a remote support connectivity option for call home and remote access functions.

rsc
Indicates the connection status (enabled or disabled) of the remote support center service.

- Enabled - the remote support center service is started.
- Disabled - the remote support center service is stopped.

If the client is not connected to an HMC, or the HMC is not running rsc, a dash (-) indicates unknown status.

Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

manageaccess

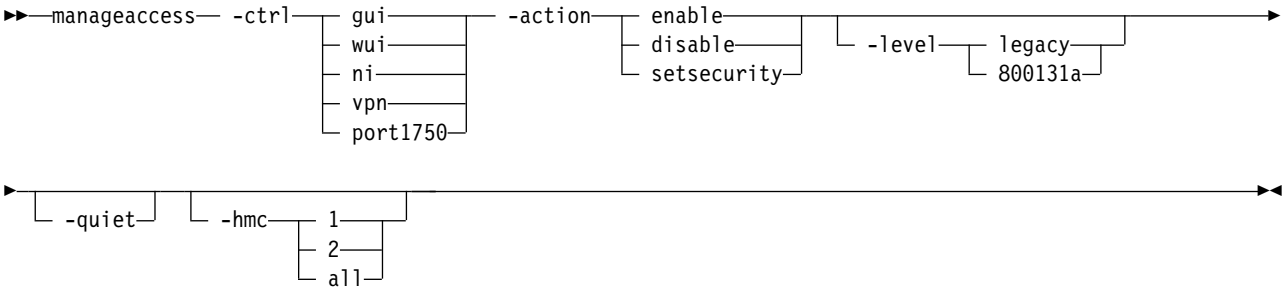
The **manageaccess** command manages the security protocol access settings of a hardware management console (HMC) for all communications to and from the DS8000 system. You can also use the **manageaccess** command to start or stop outbound virtual private network (VPN) connections instead of using the **setvpn** command.

When a DS8000 system is installed, the HMC sends a certificate (signed public key) to IBM for server authentication. A certificate is also sent for encrypting communications of applications that use VPN (Internet) connection. You can use the **manageaccess** command to start or stop a VPN session and to create a secure VPN connection. The IBM VPN server also provides extra authentication to allow traffic to certain IBM servers only, for the call home feature and for remote service.

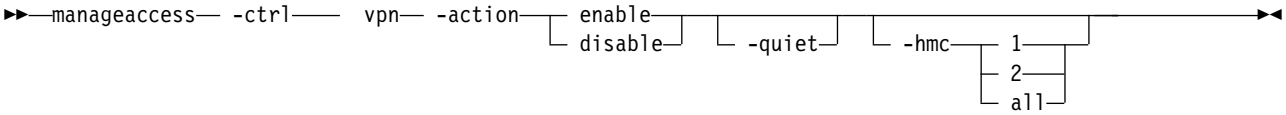
Notes:

1. Only users with admin authority can access the **manageaccess** command.
2. Only support personnel with special access rights can use the VPN connection.
3. The VPN connection is used when remote access is required by support personnel and no local onsite access to the system is available.
4. The VPN connection can take 2 - 10 minutes for the secure connection to be established and recognized by the RS3/RS4 server.
5. The secure VPN connection ends automatically when the terminal emulation session ends. However, you can end the session earlier by using the **manageaccess –action disable -ctrl vpn** command.

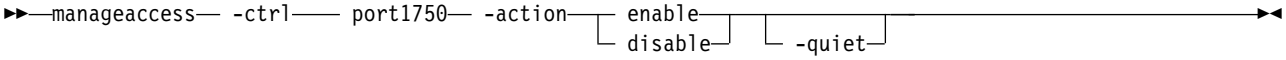
See the related link for a description of the behavior of the **–hmc** parameter.



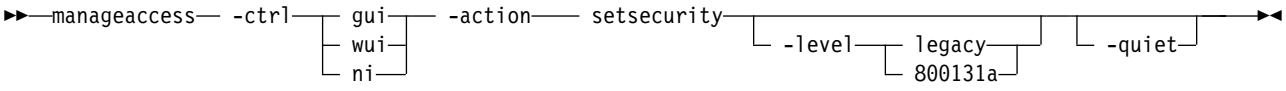
The following diagram shows the **manageaccess** command with the **-ctrl** parameter value set to **vpn**. The valid **-action** values for this control are `enable` and `disable`.



The following diagram shows the **manageaccess** command with the **-ctrl** parameter value set to **port1750**. The valid **-action** values for this control are `enable` and `disable`.



The following diagram shows the **manageaccess** command with the **-ctrl** parameter value set to **gui** | **wui** | **ni**. The valid **-action** value for these controls is `setsecurity`. You must also specify the **-level** parameter.



Parameters

-ctrl | gui | wui | ni | vpn | port1750

gui This control specifies the inbound connection protocol to the DS8000 Storage Management GUI through browsers such as Internet Explorer or Firefox. Set this control to the 800131a level only when all browsers that are accessing the DS8000 Storage Management GUI can support the NIST SP800-131a protocol.

The valid **-action** value for this control is `setsecurity`. You must also specify the **–level** value, either `legacy` or `800131a`.

For information about browser access to the Web User Interface (WUI), see the following definition.

Note: Setting this control causes the DS8000 Storage Management GUI to restart. All DS8000 Storage Management GUI users are forced to log on again.

wui This control specifies the inbound connection protocol to the service GUI or Web User Interface (WUI) through browsers such as Internet Explorer or Firefox. Set this control to the 800131a level only when all browsers that are accessing the WUI can support the NIST SP800-131a protocol.

The valid **-action** value for this control is `setsecurity`. You must also specify the **–level** value, either `legacy` or `800131a`.

For information about browser access to the GUI, see the preceding definition.

Note: Setting this control requires manually rebooting the HMC for the change to take effect.

ni This control specifies the inbound connection protocol to the DS8000 system from external programs such as the DS CLI, Copy Services Manager, and others, through the network interface (NI). Set this control to the 800131a level only when all external programs that are accessing the DS8000 system can support the NIST SP800-131a protocol.

The valid **-action** value for this control is `setsecurity`. You must also specify the **–level** value, either `legacy` or `800131a`.

Notes:

1. Setting the security mode on this control from `legacy` to `800131a` disables the port1750 control. Setting the security mode on this control from `800131a` to `legacy` restores the port1750 control to its previous value, either disabled or enabled
2. Setting this control causes the NI server to restart. Any users who meet the specified security level requirements automatically reconnects. However, all commands from the NI client programs, such as DS CLI and Copy Services Manager, fail until the NI server completes rebooting.

vpn Specifies the outbound IBM service connection. The valid **-action** values for this control are `enable` and `disable`. Setting this control does not cause any server to reboot.

Note: Starting with DS8000 Release 8.0, VPN is no longer available as a remote support connectivity option for call home and remote access functions.

port1750
Specifies the inbound port 1750 to the DS8000 system from external legacy programs, such as the DS CLI , Copy Services Manager, and others that were created before Release 7.2. Release 7.2 and later programs use the new 1751 port with an NIST SP 800-131a-compliant certificate. The 1750 port uses a certificate that is not compliant with NIST SP 800-131a.

The valid **-action** values for this control are `enable` and `disable`.

Notes:

1. If the ni control is set to `legacy`, you can use this control to enable or disable the 1750 port. If the `ni` control is set to `800131a`, this control is set to **disabled** and cannot be enabled.
2. Setting the security mode on the ni control from `legacy` to `800131a` disables this control. Setting the security mode on the ni control from `800131a` to `legacy` restores this control to its previous value, either disabled or enabled.
3. Setting this control does not cause any server to restart.

-action enable | disable | setsecurity
Specifies an action to apply to the specified access control, as follows:

enable
Specifies that the access control is enabled. This value is valid when the **-ctrl** value is `vpn` or`port1750`.

disable
Specifies that the access control is disabled. This value is valid when the **-ctrl** value is `vpn` or `port1750`.

setsecurity
Specifies the security level to use with the specified access control. This value is valid when the **-ctrl** value is , `gui`, `wui`, or `ni`.

-level legacy | 800131a
(Optional) Specifies the security protocol level. Each connection automatically uses the highest level security protocol that is supported by both ends of the connection.

This parameter is valid for the following controls: , `gui`. `wui`, and `ni`, with the **-action** value of `setsecurity`.

The following security protocols are allowed:

legacy SSLv3, TLSv1.0, TLSv1.1, TLSv1.2, and later.

800131a
TLSv1.2 and later.

Note: NIST SP 800-131a requires a minimum security level of TLSv1.2 or later.

-quiet
(Optional) Turns off the confirmation prompt for the **manageaccess** command.

-hmc 1 | 2 | all
(Optional) Specifies the primary (1), secondary (2), or both (all) HMCs for which access is modified.

Notes:

1. `-hmc 1` specifies the primary HMC, and `–hmc 2` specifies the secondary HMC regardless of how `-hmc1` and `-hmc2` were specified during the DS CLI startup. A DS CLI connection can succeed if, at DS CLI startup, a user inadvertently uses `–hmc2` to specify the primary HMC and `–hmc1` to specify the secondary HMC. Therefore, using `-hmc 1` in this command still refers to the actual primary HMC.
2. This parameter is valid only with the **-ctrl** `vpn` option and **-action** values of `enable` or `disable`. All other controls apply to all HMCs unless otherwise specified.

Example

Invoking the manageaccess command

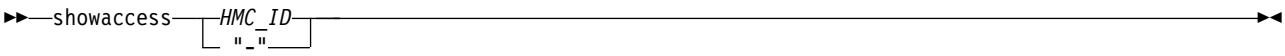
```
dscli> manageaccess –ctrl vpn –action enable
```

The resulting output

```
CMUN00022E manageaccess: The specified operation is unsupported.
```

showaccess

The **showaccess** command displays the access properties of a specified Hardware Management Console (HMC).



Parameters

HMC_ID | **-**

(Required) Specifies the primary (hmc1) or the secondary (hmc2) HMC for which the properties are displayed.

Note: The hmc1 value specifies the primary HMC, and the hmc2 value specifies the secondary HMC, regardless of how -hmc1 and -hmc2 were specified during the DS CLI startup. A DS CLI connection can succeed if you inadvertently specify a primary HMC as –hmc2 and the secondary HMC as -hmc1 when you start the DS CLI. If so, by using hmc1 in this command still refers to the actual primary HMC.

Example

Invoking the showaccess command

```
dscli> showaccess hmc1
```

The resulting output

```
dscli> showaccess hmc 1
hmc               hmc1
cmdline           enabled
wui               enabled
vpn               disabled
aos               disabled
rsc               enabled
port1750          enabled
gui(security)     800131a
ni(security)      legacy
wui(security)     legacy
```

Report field definitions

hmc Indicates the HMC for which access properties are listed. hmc1 is the primary HMC, and hmc2 is the secondary HMC.

cmdline
Indicates whether the command-line shell is enabled or disabled on the HMC through Internetor VPN connection. This information is for service access only and is not related to access to the system through the DS command-line interface. If the client is not connected to an HMC, a dash (-) indicates that the information is not available.

wui Indicates whether the Web User Interface (WUI) is enabled or disabled on the HMC through Internetor VPN connection. This information is for service access only and has no effect on access to the system by the DS8000 Storage Management GUI. If the client is not connected to an HMC, a dash (-) indicates that the information is not available.

vpn Indicates whether the connection status for the outbound VPN connection is enabled or disabled. If the client is not connected to an HMC, a dash (-) indicates that the information is not available.

Note: Starting with DS8000 Release 8.0, VPN is no longer available as a remote support connectivity option for call home and remote access functions.

aos Indicates whether the Assist On-Site (aos) connection to the HMC is enabled or disabled for IBM service. If the client is not connected to an HMC or if the query is not supported by the server, a dash (-) indicates that the information is not available.

rsc Indicates the connection status (enabled or disabled) of the remote support center service.

- Enabled - the remote support center service is started.
- Disabled - the remote support center service is stopped.

If the client is not connected to an HMC, or the HMC is not running rsc, a dash (-) indicates unknown status.

port1750
Indicates the status (enabled or disabled) of the legacy port 1750. If the client is not connected to an HMC, a dash (-) indicates that the information is not available.

gui (security)
Indicates the security protocol that is allowed for the inbound connections to the DS8000 Storage Management GUI server on the HMC or HMCs. If the client is not connected to an HMC, a dash (-) indicates that the information is not available.

The allowed security protocols are as follows:

`legacy`
SSLv3, TLSv1.0, TLSv1.1, TLSv1.2, and later.

`800131a`
TLSv1.2 and later.

ni (security)
Indicates the security protocol that is allowed for the inbound connections to the network interface (NI) server on the HMC or HMCs that the DS CLI, TPC, TPC-R, and other programs use. If the client is not connected to an HMC, a dash (-) indicates that the information is not available.

The allowed security protocols are as follows:

`legacy`
SSLv3, TLSv1.0, TLSv1.1, TLSv1.2, and later.

`800131a`
TLSv1.2 and later.

wui (security)
Indicates the security protocol that is allowed for the inbound connections to the service GUI, or WUI, server on the HMC or HMCs. The allowed security protocols are as follows:

`legacy`
SSLv3, TLSv1.0, TLSv1.1, TLSv1.2, and later.

`800131a`
TLSv1.2 and later.

Storage complex configuration commands

Various commands are available to configure a storage complex for DS8000 models.

The following storage complex configuration commands are available:

`chsp` Modifies a storage complex for items such as notification of the Simple Network Management Protocol (SNMP) traps and email problem notification lists in a storage complex.

lsvpn Displays a report that lists the outbound VPN information such as management consoles and connection status.

Note: Starting with DS8000 Release 8.0, VPN is no longer available as a remote support connectivity option for call home and remote access functions.

setvpn Used when remote access is required by IBM Support personnel and there is no local onsite access to the storage system.

Note: Starting with DS8000 Release 8.0, VPN is no longer available as a remote support connectivity option for call home and remote access functions.

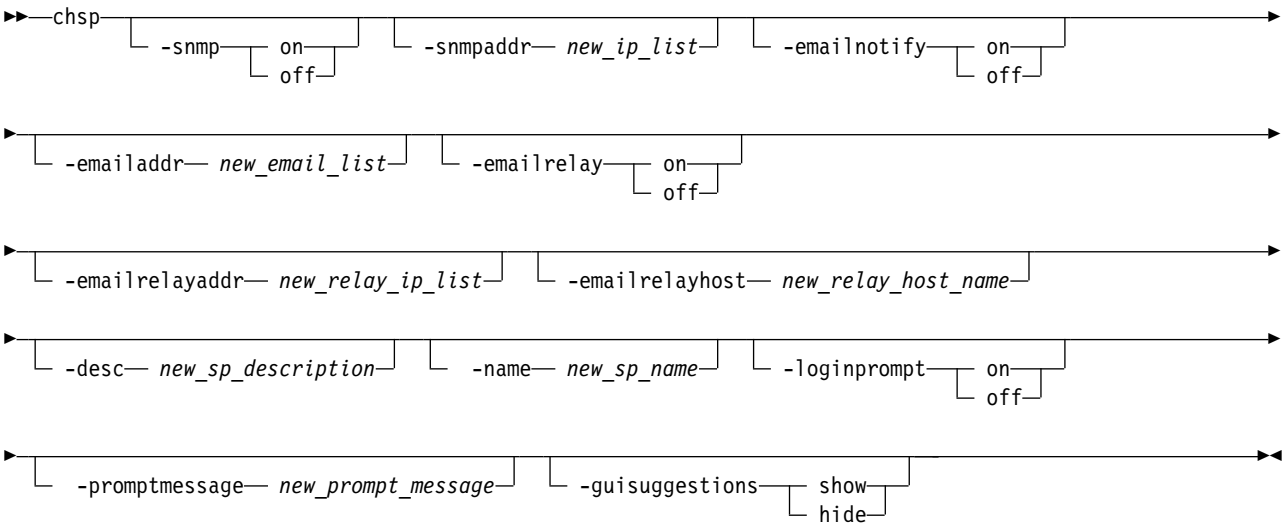
showsp Generates a report that displays detailed properties of a storage complex.

chsp

The **chsp** command modifies a storage complex for items such as notification of the Simple Network Management Protocol (SNMP) traps and email problem notification lists in a storage complex.

Notes:

- The Storage Manager server domain is a single storage complex. The storage complex object can be created or deleted only by service personnel.
- The SNMP settings apply to only some types of SNMP notifications (Example: Copy Services).



Parameters

-snmp on | off
(Optional) Specifies whether the Simple Network Management Protocol (SNMP) trap notification function sends notifications when a problem occurs on a storage complex. SNMP traps that are generated by the storage complex are sent to the IP address that is specified by the **-snmpaddr** parameter.

-snmpaddr *new_ip_list*
(Optional) Specifies a new SNMP trap destination list. This list consists of one or two IP addresses that receive SNMP traps that are generated by the storage complex if SNMP is enabled.

Note: Multiple IP addresses must be separated with a comma with no space before or after each comma.

-emailnotify on | off
(Optional) Specifies whether email notifications are sent to the designated email addresses when a problem occurs on a storage complex.

When you enable the **-emailnotify** and **-emailrelay** parameters, the email notification is sent directly to the IP address that is associated with the **-emailrelayhost** *hostname* parameter. However, if the email relay host name is not specified, the email notification is sent directly to the email relay address.When email notify is enabled and email relay is disabled, the email notification is sent directly to the specified email address.

-emailaddr *new_email_list*
(Optional) Specifies the email IP address or addresses where problem notifications are sent. You can specify two IP addresses.

Note: Multiple IP addresses must be separated with a comma with no space before or after the comma.

-emailrelay on | off
(Optional) When you enable the **-emailrelay** and **-emailnotify** parameters, the email notification is sent directly to the IP address that is associated with the **-emailrelayhost** *new_relay_host_name* parameter. However, if the email relay host name is not specified, the email notification is sent directly to the email relay address.

When email notify is enabled and email relay is disabled, the email notification is sent directly to the specified email address.

-emailrelayaddr *new_relay_ip_list*
(Optional) Specifies the email relay IP address list. You can specify two IP addresses.

Note: Multiple IP addresses must be separated with a comma with no space before or after the comma.

-emailrelayhost *new_relay_host_name*
(Optional) Specifies the new email relay host system name.

-desc *new_sp_description*
(Optional) Specifies your description of the storage complex. This description is limited to 256 byte or 128 double-byte characters.

-name *new_sp_name*
(Optional) Specifies the name that you designate for the storage complex. This name is limited to 32 byte or 16 double-byte characters.

-loginprompt on | off
(Optional) Specifies whether the login prompt is enabled or disabled.

-promptmessage *new_prompt_message*
(Optional) Specifies the detailed prompt message. You cannot modify prompt message if the -loginprompt parameter is set to `off`.

-guisuggestions show | hide
(Optional) Specifies whether suggested tasks are shown or hidden in the task icon at the top right of the DS8000 Storage Manager GUI. The default is `show`.

Invoking the `chsp` command

```
dscli> chsp -desc "my storage complex"
```

The resulting output

Storage-complex IBM.2107-75FA120 successfully modified.

setvpn

The **setvpn** command starts or ends an outbound virtual private network connection (VPN).

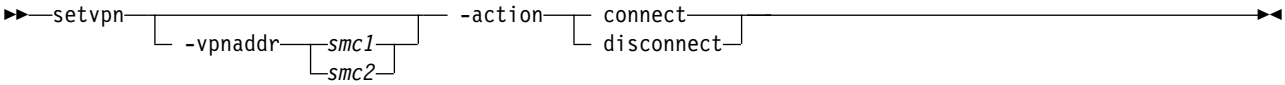
Note: VPN is no longer supported as a remote support connectivity option for call home and remote access functions. As a result, starting with DS8000 Release 8, the **setvpn** command is no longer supported.

During the installation of the DS8000 product, the hardware management console (HMC) sends a certificate (signed public key) to IBM for server authentication and for SSL encryption of applications using VPN (Internet and modem) connections.

You can use the **setvpn** command to start or stop the session and to create a secure connection. In addition, the IBM VPN server does additional authentication to allow traffic to certain IBM servers only, for the call home feature and remote service.

Notes:

1. Only technical support personnel with special access rights can use the VPN connection.
2. The **setvpn** command is used when remote access is required by technical support personnel and there is no local on-site access to the machine.
3. It can take from 2 to 10 minutes for the secure connection to be established and recognized by the RS3/RS4 server.
4. The secure connection ends automatically when the terminal emulation session ends. However, you also have the ability to end the session earlier by issuing the **setvpn -action** *disconnect* command.
5. The **-vpnaddr** parameter requires that you specify a value for either *smc1* or *smc2*. If you do not specify the **-vpnaddr** parameter, the storage management console (SMC) for the current connection is used. The SMC address is taken from the profile file or the SMC address that you specify on the DS CLI command line.



Parameters

-vpnaddr *smc1* | *smc2*
(Optional) Specifies the VPN server machine. In addition, you can specify where you want the outbound VPN to start from by designating the following values:

smc1 Identifies the management console (SMC) where you want the outbound VPN to start from. The console that you have specified in your profile for *hmc1* starts your DS CLI session, unless you specify a console that is not designated in your profile. In this case, the console that you specify to start your session is the one where the connection is made.

smc2 Identifies the management console where you want the outbound VPN to start from. The console that you have specified in your profile for *hmc2* starts your DS CLI session, unless you specify a console that is not designated in your profile. In this case, the console that you specify to start your session is the one where the connection is made.

-action connect | disconnect
(Required) Specifies that the secure VPN connection be started or disconnected.

Example

Invoking the setvpn command

```
dscli> setvpn –vpnaddr smc1 –action connect
```

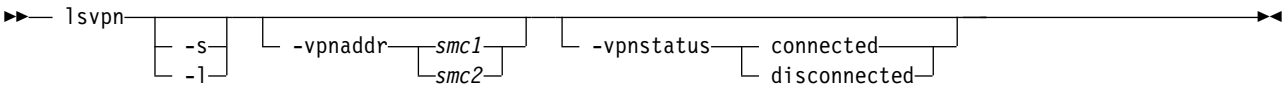
The resulting output

```
Sun July 10 02:23:49 PST 2016 DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
Secure connection started.
```

lsvpn

The **lsvpn** command displays a report that lists the outbound VPN information such as management consoles and connection status.

Note: VPN is no longer supported as a remote support connectivity option for call home and remote access functions. As a result, starting with DS8000 Release 8, the **lsvpn** command is no longer supported.



Parameters

-s
(Optional) Displays the management console (SMC) and connection status. You cannot use the **-l** and **-s** parameters together.

-l
(Optional) Displays the management console (SMC) and connection status. You cannot use the **-l** and **-s** parameters together.

-vpnaddr *smc1* | *smc2*
(Optional) Specifies the VPN server machine. In addition, you can specify where the outbound VPN was started from by designating the following values:

smc1 Identifies the management console (SMC) where the outbound VPN was started from. The console in your profile for hmc1 starts your DS CLI session, unless you specify a console that is not designated in your profile. In this case, the console you specified to start your session is the one reported where the connection was made.

smc2 Identifies the management console where the outbound VPN was started from if you did not start it from the management console that is identified by *smc1*. The console in your profile for hmc2 starts your DS CLI session, unless you specified a console that is not designated in your profile. In this case, the console that you specified to start your session is the one reported where the connection was made.

Note: The default value is to display the address for *smc1* and *smc2*. If you do not specify the **-vpnaddr** parameter, the generated report displays both the *smc1* and *smc2* addresses as they are recorded in your profile.

-vpnstatus *connected* | *disconnected*
(Optional) Specifies that you receive a report that displays only the SMC for the connection status specified.

Note: The default value is to display the connection status for all SMCs. The generated report displays all connected and disconnected SMCs.

Example

For this command and all other DS CLI list commands, the results are shown in table format for clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lsvpn** command.

Invoking the lsvpn command

```
dscli> lsvpn -l
```

The resulting output

VPN Address	VPN Status
smc1	Connected
smc2	Disconnected

showsp

The **showsp** command displays a properties report for a storage complex.

The report included the properties values for the names, descriptions, and customer account names of the storage complex.



Parameters

No parameters exist for this command.

Notes:

1. The SNMP settings apply to only some types of SNMP notifications (Example: Copy Services).
2. The email parameters are not currently supported.

Example

The following table represents the headers that are displayed on the output report that is associated with the **showsp** command.

Invoking the showsp command

```
dscli> showsp
```

The resulting output

Name	Desc	Acct	SNMP	SNMPadd
My_ storage-complex	Production storage-complex	ABC Company	Enabled	9.*xxx*.14.245

eMailnotify	eMailaddr	eMailrelay	eMailrelay-addr	eMailrelay-host	numks supported
Enabled	email1@ ibmds8000.com, email2@ ibmds1.com	Disabled	9.*xxx*.14.45	relay_host	1

loginprompt	promptmessage	guisuggestions
On	ese switch is on	Show

Report field definitions

Name Indicates the name of the storage complex.

Desc Indicates the description of the storage complex.

Acct Indicates the customer account name for the storage complex.

SNMP
Indicates whether SNMP traps that are generated by the storage complex for software events are sent to the IP address that is specified by the **chsp** command. Other types of SNMP traps are enabled or disabled by using the user interface on the HMC. This column displays the values: Enabled or Disabled.

SNMPadd
Indicates one or two IP addresses if SNMP is enabled. These addresses indicate where SNMP traps that are generated by the storage complex are sent. Multiple IP addresses are separated with commas with no blank space before and after each. For example: 9.111.14.254,9.113.22.236

eMailnotify
Indicates whether email notifications are enabled or disabled.

When the **chsp** command enables the **-emailnotify** and **-emailrelay** parameters, the email notification is directed to the IP address that is associated with the **-emailrelayhost** *hostname* parameter. However, if the email relay host name is not specified, the email notification is directed to the email relay address.

When email notify is enabled and email relay is disabled, the email notification is sent directly to the specified email address.

eMailaddr
Indicates one or more email addresses to which notification is sent if service is required when email is enabled. A dash (-) is displayed if email is not enabled or if an email address is not available.

eMailrelay
Indicates whether **-emailrelay** is enabled or disabled.

When **-email** and **-emailrelay** parameters are enabled, the email notification is directed to the IP address that is associated with the **-emailrelayhost** *hostname* parameter. However, if the email relay host name is not specified, the email notification is directed to the email relay address.

When **-email** is enabled and **-emailrelay** is disabled, the email is sent directly to the specified email address.

eMailrelayaddr
Indicates one or more email relay IP addresses. These IP addresses represent the addresses through which notification is relayed if service is required when email is enabled.

eMailrelayhost
Indicates the email relay host system name.

numkssupported
Indicates the number of key servers that are supported.

loginprompt
Indicates the following values of the login message prompt:

On
Indicates that the login message prompt is enabled.

Off
Indicates that the login message prompt is disabled.

"-"
Indicates that the message prompt is not supported.

promptmessage
Indicates the login message prompt. A dash (-) displays if the promptmessage option is not supported.

guisuggestions
Indicates the following values for the GUI suggestions:

Show
Indicates suggested tasks display in the task icon at the top right of the DS8000 Storage Management GUI.

Hide
Indicates suggested tasks do not display in the task icon at the top right of the DS8000 Storage Management GUI.

Storage configuration commands

There are some specific DS CLI commands that are associated with configuring storage for IBM Z and open system hosts.

The commands below allow you to configure storage for IBM Z and open system hosts.

Array site specific commands

Specific commands are used to display array site information.

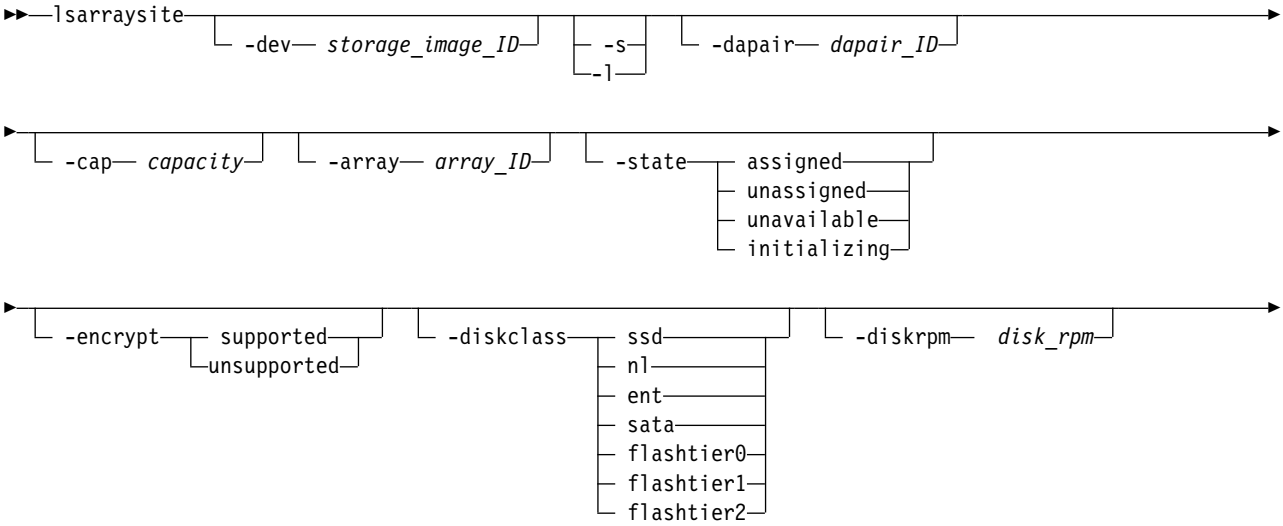
The following array site commands are available:

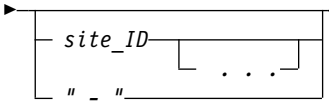
lsarraysite
Generates a report that lists the array sites and status information for each array site in the list.

showarraysite
Generates a report that displays the detailed properties of a specific storage image array site.

lsarraysite

The **lsarraysite** command displays a list of array sites and status information for each array site in the list.





Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified site ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-s
(Optional) Displays the array ID. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output plus the disk drive module rpm (revolutions per minute), disk class, and encryption support capability. You cannot use the **-l** and the **-s** parameters together.

-dapair *dapair_ID*
(Optional) Displays array sites that are associated with a common device adapter pair ID. A device adapter pair ID is a two-digit decimal number with no leading zeros.

-cap *capacity*
(Optional) Displays in decimal gigabytes (GB) the array sites that have the specified capacity of the disk drive module.

-array *array_ID*
(Optional) Displays the array site that is associated with the specified array ID. An array ID is a four-digit decimal number with no leading zeros, prefixed with the letter *A*.

-state assigned | unassigned | unavailable | initializing
(Optional) Displays array sites that are in the specified state. One of the following values is displayed:

assigned
Specifies that the designated array site is defined as an array.

unassigned
Specifies that the array site is available to be defined as an array.

unavailable
Specifies that the designated array site is unassigned and at least one disk is not in the normal state. Also, the array site is not in the initializing state.

initializing
Specifies that the array site is unassigned and all disks are either in the normal or initializing state. Also, at least one disk is in the initializing state.

-encrypt supported|unsupported
(Optional) Displays only the array that has the specified encryption capability.

-diskclass ssd | nl | ent | sata | flashtier0 | flashtier1 | flashtier2
(Optional) Displays only array sites with the following specified disk classes:

ssd Specifies solid-state devices.

nl Specifies high-capacity near-line disk drives.

ent Specifies high-speed Enterprise disk drives.

sata Specifies high-capacity SATA devices.

flashtier0
Specifies high performance flash (Flash Tier 0) drives.

flashtier1
Specifies the first tier of high capacity flash (Flash Tier 1) drives.

flashtier2
Specifies the second tier of high capacity flash (Flash Tier 2) drives.

-diskrpm *disk_rpm*
Displays only array sites with the specified disk RPM.

site_ID . . . | **-**
(Optional) Displays array sites that have the specified IDs. An array site identifier is a four-digit decimal number with no leading zeros, prefixed with the letter *S*.

To specify a range of array site IDs, separate the array site IDs with a hyphen.

You must separate multiple array site IDs or ranges of array site IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lsarraysite** command using the **-l** parameter.

Invoking the lsarraysite command

```
dscli> lsarraysite -l
```

The resulting output

Arsite	DA pair	Dkcap (10^9B)	Diskrpm	State	Array	Diskclass	Encrypt
S1	2	146.0	10000	Assigned	A1	ENT	supported
S2	2	146.0	10000	Assigned	A0	ENT	supported
S3	2	146.0	10000	Assigned	A3	NL	unsupported
S4	2	146.0	10000	Assigned	A2	SATA	unsupported

Report field definitions

Arsite*
Indicates the array site ID. The array site ID is a four-digit decimal number, with no leading zeros, prefixed with the letter S.

Note: The array site ID does not point to a physical location.

DA pair
Indicates the DA pair ID. DA pairs are located in I/O enclosure pairs. The DA pair ID indicates the I/O enclosure location.

Note: An even-numbered DA pair ID indicates the first DA pair in an I/O enclosure pair. An odd-numbered DA pair ID indicates the second DA pair in an I/O enclosure pair.

Dkcap (10^9B)
Indicates the minimum disk capacity of the disks in the designated array site in decimal gigabytes (GB).

Diskrpm+
Indicates the minimum disk rpm of the disks in the designated array site.

State
Indicates the array site state. One of the following values can be displayed in this field:

Assigned
Indicates that the designated array site is defined as an array.

Unassigned
Indicates that the array site is available to be defined as an array.

Unavailable
Indicates that the designated array site is unassigned and at least one disk is not in the normal state. Also, the array site is not in the initializing state.

Initializing
Indicates that the array site is unassigned and all disks are either in the normal or initializing state. Also, at least one disk is in the initializing state.

Array
Indicates the array ID that this assigned array site is assigned to. The ID is prefixed by the letter *A*.

Diskclass+
Indicates the disk class. One of the following values can be displayed:

ENT
Indicates enterprise and represents high-speed Fibre Channel disk drives.

NL Indicates near-line and represents ATA (FATA) disk drives

SATA
Indicates high capacity SATA disk drives.

SSD
Indicates solid-state devices.

FlashTier0
Indicates high performance flash (Flash Tier 0) drives.

FlashTier1
Indicates the first tier of high capacity flash (Flash Tier 1) drives.

FlashTier2
Indicates the second tier of high capacity flash (Flash Tier 2) drives.

encrypt+
Indicates the encryption support capability. One of the following values can be displayed:

supported
The disk drive modules in this arraysite are capable of encryption.

unsupported
The disk drive modules in this arraysite are not capable of encryption.

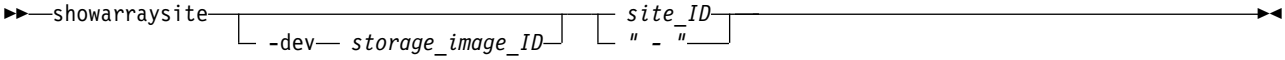
Key:

* Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

showarraysite

The **showarraysite** command displays detailed properties of a specific storage image array site.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified site ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

site_ID | **-**
(Required) Specifies that information be displayed for the designated array site ID. This parameter also accepts a fully qualified site ID, which consists of the storage image ID or a shortened version without the storage image ID, if the **-dev** parameter is specified. The shortened version is a four-digit decimal number with no leading zeros, prefixed by the letter *S*. The array site ID does not imply a physical location.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **showarraysite** command.

Invoking the showarraysite command

```
dscli> showarraysite -dev IBM.2107-75FA120 S11
```

The resulting output

Arsite	DA pair	Dkcap (10^9B)	Diskrpm	State
IBM.2107-75FA120/S11	IBM.2107-75FA120/0	146	15000	Assigned

Array	Dkinf	Dkrate (GB/sec)	DDMSN	Spares	dataDDM	Diskclass	Encrypt
IBM.2107-75FA120/A44	FCAL	2	0123456789 ABCDEF	0	8	SATA	Unsupported

Report field definitions

Arsite Indicates the array site ID. The array site ID is a decimal number up to four digits in length, no leading zeros, prefixed by the letter *S*. For example, S11.

Notes:

1. Array sites consist of eight DDMs.
2. The array site ID does not imply a physical location.

DA pair
Indicates the DA pair ID. DA pairs are located in I/O enclosure pairs. DA pair ID implies the I/O enclosure location.

Note: DA Adapters are installed in slot 3 in one enclosure and slot 6 in the peer enclosure. The DA pair ID identifies the enclosure that contains the DA Adapter in slot 3. For example, a DA adapter is installed in slot of 3 of enclosure 2. Its peer is installed in slot 6 of enclosure 3. In this case, the DA Pair ID is 2.

Dkcap (10^9B)
Indicates the minimum disk capacity of the disks in the designated array site.

Diskrpm
Indicates the minimum disk rpm of the disks in the designated array site.

State Indicates the array site state. The values that can be displayed in this field are as follows:

assigned
Indicates that the designated array site is defined as an array.

unassigned
Indicates that the array site is available to be defined as an array.

Array Indicates the array ID that the designated array site is assigned to. The ID is prefixed by the letter *A*.

Dkinf Indicates the DDM interface type for the disks in this array site. The following value is displayed:
- FC-AL (Fibre Channel arbitrated loop)
- SAS (Serial Attached SCSI)

Dkrate
Indicates the minimum disk interface rate of the disks in the designated array site.

DDMSN
Indicates the list of DDM serial numbers (SN) that are associated with the designated array site. Each DDM SN is a 16-character string. Each serial number is separated by a comma.

Spares
Indicates the number of spare DDMs that are allocated from the array site.

DataDDM
Indicates the number of data DDMs. This value is based on the number of DDMs minus the number of spares.

Diskclass
Indicates the disk class. One of the following values is displayed:

ENT Indicates enterprise and designates a high-speed Fibre Channel disk drive.

NL Indicates near-line and represents ATA (FATA) disk drives.

SATA Indicates high-capacity SATA disk drives.

SSD Indicates solid-state devices.

FlashTier0
Indicates high performance flash (Flash Tier 0) drives.

FlashTier1
Indicates the first tier of high capacity flash (Flash Tier 1) drives.

FlashTier2
Indicates the second tier of high capacity flash (Flash Tier 2) drives.

encrypt
Indicates the encryption support capability. One of the following values are displayed:

supported
The disk drive modules in this array site are capable of encryption.

unsupported
The disk drive modules in this array site are not capable of encryption.

Array specific commands

Array specific commands are used to create and delete arrays and to display array information.

The following array specific commands are available:

lsarray
Generates a report that displays a list of arrays in a storage image and the status information for each array in the list.

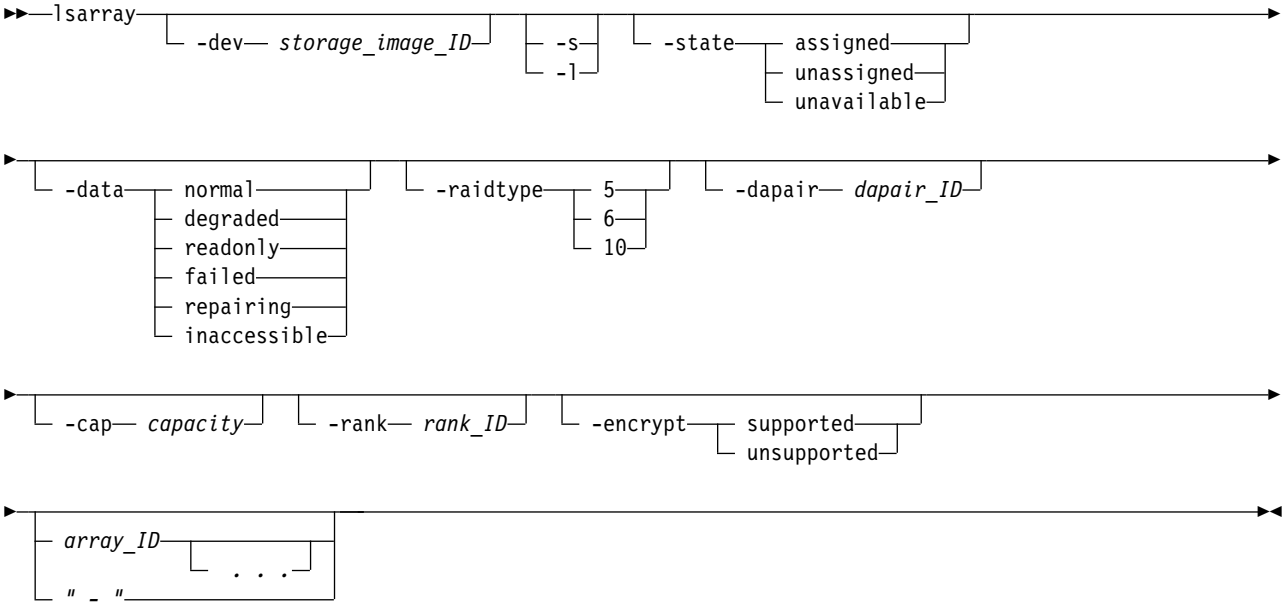
mkarray
Creates one array per command.

rmarray
Removes the specified array or arrays from the storage system.

showarray
Generates a report that displays the detailed properties of a specific array.

lsarray

The **lsarray** command displays a list of arrays in a storage image and status information for each array in the list.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes the manufacturer, machine type, and serial

number. The storage image ID is required if you do not specify a fully qualified array ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-s
(Optional) Displays only the array ID. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays default output plus the disk class and encryption capability of the arrays. You cannot use the **-l** and the **-s** parameters together.

-state assigned | **unassigned** | **unavailable**
(Optional) Specifies that information about those arrays in the designated state are to be displayed on the generated report.

-data normal | **degraded** | **readonly** | **failed** | **repairing** | **inaccessible**
(Optional) Specifies that information about those arrays in the designated data state are to be displayed on the generated report.

-raidtype 5 | **6** | **10**
(Optional) Displays only those arrays with the specified RAID type, 5, 6, or 10.

-dapair *dapair_ID*
(Optional) Displays only the array that is specified by the device adapter pair ID. A device adapter pair ID is a two-digit decimal number with no leading zeros.

-cap *capacity*
(Optional) Displays in decimal gigabytes (GB) only the array with the specified DDM capacity. You can specify up to three digits after the decimal point, for example -cap 1.447.

-rank *rank_ID*
(Optional) Displays only the array that is assigned to the specified rank ID. A rank ID is a four-digit decimal number with no leading zeros, prefixed with the letter *R*.

-encrypt supported | **unsupported**
(Optional) Displays only the array sites that have the specified encryption capability.

array_ID . . . | **-**
(Optional) Displays array information for the specified arrays. An array ID is a four-digit decimal number with no leading zeros, prefixed with the letter *A*.

To specify a range of array IDs, separate the array IDs with a hyphen. For example: A10-A12 (equates to A10 A11 A12)

You must separate multiple array IDs or ranges of array IDs with a blank space between each ID or range of IDs. For example: A11 A12 A14-A16. Your command in this case could look similar to this example:

```
dscli> lsarray IBM.2107-75FA120 -l A11 A12 A14-A16
```

The ellipsis (...) indicates that, optionally, you can specify multiple values.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lsarray** command using the **-l** parameter.

Invoking the lsarray command

```
dscli> lsarray -dev IBM.2107-75FA120 -l
```

The resulting output

Array	State	Data	RaidType
A10	Assigned	Normal	5(6+p)
A11	Assigned	Normal	5(7+p)
A12	Assigned	Normal	5(6+p)
A13	Unassigned	Normal	5(7+p)

Report field descriptions

Array*

Indicates the array number. The array number starts with the prefix *A*, followed by a decimal number up to four digits in length, with no leading zeros. For example, A44.

State

Indicates the relationship between the array and a rank. One of the following values is displayed:

Assigned

The array is assigned to a rank.

Unassigned

The array is not assigned to a rank and all of the storage devices that are indicated by the disk serial numbers attribute are in the normal state.

Unavailable

The array is not assigned to a rank and one or more of the disk drive modules (DDMs) that are indicated by the disk serial numbers attribute are not in the normal state.

Data

This value reflects the current data status. One of the following values is displayed:

Normal

The array is in the Normal data state if none of the other data states applies. This status applies if the array is unassigned.

Degraded

The array is in the Degraded data state if both of the following conditions exist:

- The Read-only, Failed, Repairing, or Inaccessible data states do not apply.
- One or more redundancy groups are rebuilding (that is, there is a DDM with a rebuilding state in the array).

Read Only

The array is in the Read-only data state if all of the following conditions exist:

- The Failed, Repairing, and Inaccessible data states do not apply.
- One or more DDMs have failed.
- There are insufficient spares to support all rebuild operations.
- Continued write operation without redundancy could result in data loss.

Failed

The array is in the Failed data state if all of the following conditions exist:

- The Repairing and Inaccessible data states do not apply.
- Two or more DDMs in the array have failed.
- There are insufficient DDMs left in the array to rebuild the data that was lost on the failing storage devices.

Repairing
The array is in the Repairing data state if all of the following conditions exist:

- The Inaccessible data status does not apply.
- The array has previously entered the failed state.
- The repair array function has been accepted.
- The repair array function has not completed.

Inaccessible
The array is in the Inaccessible data state if the storage system cannot access a set of storage devices that are sufficient to access all the data on the array.

RaidType
Indicates the type of RAID array (5, 6, or 10) and the array configuration (for example, 6+P). DS8000 array configurations are based on 8 DDM array sites for DA Pair Types 1 and 2, and 7 or 8 DDM array sites for DA Pair Type 3.

Note: The RAID type 6 is displayed in DS8000 models only.

Arsite
Indicates the array sites that are associated with the array.

Rank
Indicates the rank the array is assigned to. The value is displayed as a combination of a Storage Image ID and a rank number. The rank number is the prefix *R*, followed by a decimal number up to four digits in length, with no leading zeros. For example, R26.

DA pair
Identifies the DA pair ID. DA pairs are located in I/O enclosure pairs. DA pair ID indicates the I/O enclosure location.

Note: An even-numbered DA pair ID indicates the first DA pair in an I/O enclosure pair. An odd-numbered DA pair ID indicates the second DA pair in an I/O enclosure pair.

DDMcap (10^9B)
Indicates the minimum disk capacity in decimal gigabytes (GB) of the storage devices (DDMs) in the specified array.

Diskclass+
Indicates the disk class. One of the following values is displayed:

ENT Indicates enterprise and represents high speed Fibre Channel disk drives.

NL Indicates near-line and represents ATA (FATA) disk drives.

SATA Indicates high capacity SATA disk drives.

SSD Indicates solid-state devices.

FlashTier0
Indicates high performance flash (Flash Tier 0) drives.

FlashTier1
Indicates the first tier of high capacity flash (Flash Tier 1) drives.

FlashTier2
Indicates the second tier of high capacity flash (Flash Tier 2) drives.

encrypt
Indicates the encryption support capability. One of the following values is displayed:

supported
The disk drive modules in this array are encryption capable.

unsupported
The disk drive modules in this array are not encryption capable.

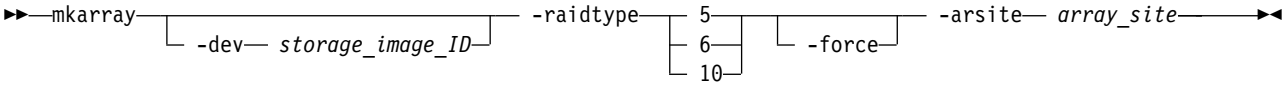
Key:

* Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

mkarray

The **mkarray** command creates one array per command. You can specify only one array site for a DS8000 machine type.



Parameters

Note: The DS8000 system assigns the ID during array creation based on the current configuration, past configuration changes, and other internal considerations. No algorithm is available to accurately predict the newly created array ID in every case.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified array site, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-raidtype 5 | 6 | 10
(Required) Specifies a RAID type for the array.

Note: The **-raidtype 6** parameter can be specified for DS8000 models only.

-force
(Optional) The force option applies only when you create a RAID 5 array. An information message displays a disclaimer about using RAID 5. If you agree with the disclaimer, the RAID 5 array creation is processed. Even when this option is not used, a disclaimer displays and you must accept the disclaimer to create RAID 5.

-arsite *array_site*
(Required for a DS8000 machine type) Specifies the array site for the array. An array site number is a four-digit decimal number with no leading zeros, prefixed with the letter *S*.

Example of fully qualified array site: IBM.2107-75FA120/S11

Specify 1 or 2 array sites for 1750 RAID types 5 and 10. If there are two array sites, both must be associated with a common DA pair ID. An array site number is a four-digit decimal number with no leading zeros, prefixed with the letter *S*. Separate the two array sites by a comma with no blank space in between. Example: S10,S11.

Example

Invoking the mkarray command

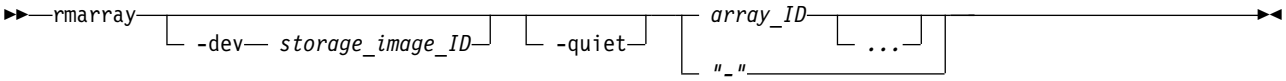
```
dscli> mkarray -dev IBM.2107-75FA120 -raidtype 10 -arsite S08
```

The resulting output

```
Array A10 successfully created.
```

rmarray

The **rmarray** command deletes arrays.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for all array IDs, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-quiet
(Optional) Turns off the array removal confirmation prompt for this command.

array_ID ... | -
(Required) Specifies the array IDs that are to be deleted. Accepts a fully qualified array ID, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is a four digit decimal number with no leading zeros, prefixed by the letter "A".

To specify a range of array IDs, separate the array IDs with a hyphen.

You must separate multiple array IDs or ranges of array IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input.

Note: You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmarray command

```
dscli> rmarray -dev IBM.2107-75FA120 A44
```

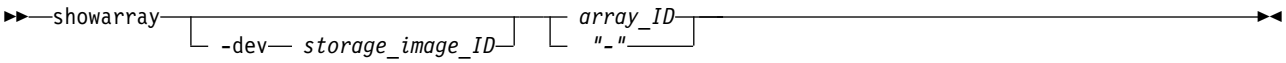
The resulting output

```
Are you sure you want to delete array IBM.2107-75FA120/A44? [y/n]: Y

Array Storage Image ID/A44 successfully deleted.
```

showarray

The **showarray** command displays detailed properties of a specific array.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified array ID. It is also

required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

array_ID | **-**
(Required) Specifies the array ID that you want to view. This parameter accepts a fully qualified array ID, which consists of the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is a four-digit decimal number with no leading zeros, prefixed by the letter *A*.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **showarray** command.

Invoking the showarray command

```
dscli> showarray -dev IBM.2107-75FA120 A44
```

The resulting output

Array	SN	State	Datastate	RaidType	Arsite
IBM.2107-75FA120/ A44	AZ123AQ	Assigned	Normal	5 (6+P)	S21

Rank	DA Pair	DDMcap (10^9B)	DDMRPM	Interface Type	Interrate (GB/secs)	Diskclass	Encrypt
2107-75FA 123/R26	2107-75FA 123/11	145	15000	FCAL	2	SATA	unsupported

Report field definitions

Array Indicates the array ID number. The array ID number starts with the prefix *A*, followed by a decimal number up to 4 digits in length, with no leading zeros. For example, A44.

SN Indicates the unique internal identifier for the data space of the designated array (for example, AZ123AQ).

State Indicates the array state. One of the following values is displayed:

Assigned
The array is assigned to a rank.

Unassigned
The array is not assigned to a rank and all of the storage devices that are indicated by the disk serial numbers attribute are in the normal state.

Unavailable
The array is not assigned to a rank and one or more of the disk drive modules (DDMs) that are indicated by the disk serial numbers attribute are not in the normal state.

Datastate
Indicates the current data state. One of the following values is displayed:

Normal
The array is in the Normal data state if none of the other data states applies. This status applies if the array is unassigned.

Degraded
The array is in the Degraded data state if both of the following conditions exist:

- The Read-only, Failed, Repairing, or Inaccessible data states do not apply.
- One or more redundancy groups are rebuilding (that is, there is a DDM with a rebuilding state in the array).

Read Only
The array is in the Read-only data state if all of the following conditions exist:

- The Failed, Repairing, and Inaccessible data states do not apply.
- One or more DDMs failed.
- There are insufficient spares to support all rebuild operations.
- Continued write operation without redundancy might result in data loss.

Failed
The array is in the Failed data state if all of the following conditions exist:

- The Repairing and Inaccessible data states do not apply.
- Two or more DDMs in the array failed.
- There are insufficient DDMs left in the array to rebuild the data that was lost on the failing storage devices.

Repairing
The array is in the Repairing data state if all of the following conditions exist:

- The inaccessible data state does not apply.
- The array previously entered the failed state.
- The repair array function was accepted.
- The repair array function did not complete.

Inaccessible
The array is in the inaccessible data state if the storage unit cannot access a set of storage devices that are sufficient to access all the data on the array.

RaidType
Indicates the type of RAID array (5, 6, or 10) and the array configuration (for example, 6+P). DS8000 array configurations are based on 8 DDM array sites for DA Pair Types 1 and 2, and 7 or 8 DDM array sites for DA Pair Type 3.

Arsite Indicates the array sites that are associated with the array.

Rank Indicates the rank that the array is assigned to. The value is displayed as a combination of a storage image ID and a rank number. The rank number is the prefix *R*, followed by a number up to 4 digits in length, with no leading zeros. For example, R26.

Note: If the array is unassigned, the field is " – "

DA pair
Indicates the DA pair ID. DA pairs are in I/O enclosure pairs. The DA pair ID indicates the location of the I/O enclosure.

Note: DA adapters are installed in slot 3 an enclosure and slot 6 in the peer enclosure. The DA pair ID identifies the enclosure that contains the DA adapter in slot 3. For example, a DA adapter is installed in slot of 3 of enclosure 2. Its peer is installed in slot 6 of enclosure 3. In this case, the DA Pair ID is 2.

DDMcap (10^9B)
Indicates the minimum disk capacity (10^9B) of the storage devices (DDMs) in the designated array.

DDMRPM
Indicates the minimum disk rpm of the DDMs in the designated array.

Interface Type
Indicates the disk interface type of the DDMs in the designated array.

Interrate
Indicates the minimum disk interface rate of the disks in the designated array.

Diskclass
Indicates the disk class. One of the following values can be displayed:

ENT Indicates enterprise and represents high-speed Fibre Channel disk drives.

NL Indicates near-line and represents ATA (FATA) disk drives.

SATA Indicates high capacity SATA disk drives.

SSD Indicates solid-state devices.

FlashTier0
Indicates high performance flash (Flash Tier 0) drives.

FlashTier1
Indicates the first tier of high capacity flash (Flash Tier 1) drives.

FlashTier2
Indicates the second tier of high capacity flash (Flash Tier 2) drives.

encrypt
Indicates the encryption support capability. One of the following values can be displayed:

supported
The disk drive modules in this array support encryption.

unsupported
The disk drive modules in this array do not support encryption.

Rank specific commands

There are specific DS CLI commands that are used to create, modify, and delete ranks and to display rank information.

The following rank specific commands are available:

chrank Assigns an unassigned rank to an extent pool or removes an assigned rank from a extent pool. This command can also be used to change an assigned rank to an unassigned rank.

lsrank Generates a report that displays a list of defined ranks in a storage system and the status information for each rank in the list.

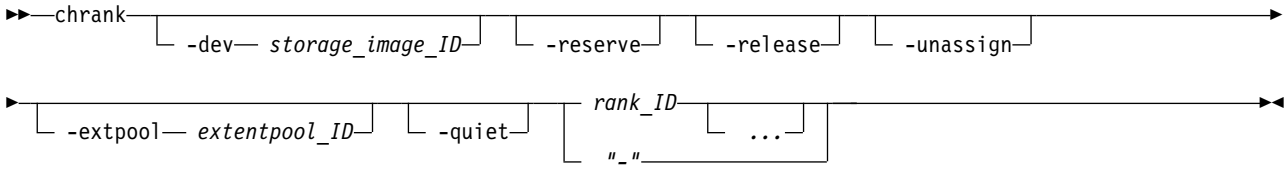
mkrank Creates one fixed block or count key data (CKD) rank from one array.

rmrank Deletes the specified ranks from a storage system.

showrank
Generates two types of reports. One report displays the detailed properties of a specified rank. The other report displays the performance metrics of a specified rank.

chrank

The **chrank** command assigns an unassigned rank to an extent pool, or removes an assigned rank from an extent pool. This command can also be used to change an assigned rank to an unassigned rank.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-reserve
(Optional) Changes the rank configuration state to *Reserved* if the rank is assigned to an extent pool, or changes the state to Unassigned Reserved if the rank is not assigned to an extent pool. If the rank is in the *depopulating* state, changing the configuration state to *Reserved* with the *-reserve* parameter stops the rank depopulation. In this case, data is not migrated back to the depopulation rank. Use the **-release** parameter to convert it back to the normal state and restart the depopulation later. The depopulation process starts from where it left off when the abort was triggered.

Note: Versions of the DS CLI before Release 6.1 fail if you use the **-reserve** parameter on Release 6.1 or later DS8000 storage systems when the rank is in the *unassigned* state. These DS CLI versions cannot handle the *Unassigned Reserved* state.

-release
(Optional) Changes the rank configuration state from Unassigned Reserved to Unassigned, or from Reserved, Normal, or Depopulation Error to Normal.

-unassign
(Optional) Changes the rank configuration state to Unassigned by removing the rank from the extent pool. If extents exist on this rank are used to provision a volume, the unassign process normally fails. If the Easy Tier LIC feature is active, then the rank configuration state is set to Depopulating while the extents are migrated to other ranks in the same extent pool. If the depopulation is successful, the rank configuration state is set to Unassigned. If it is not successful, the state is set to Depopulation Error.

Notes:

1. If the rank enters the *Depopulating* configuration state, the time that is required to depopulate the rank depends on the number of allocated extents; it might take a long time before the rank enters the *Unassigned* state. When this situation occurs, a confirmation prompt is given, but can be suppressed with the **-quiet** parameter.
2. Versions of the DS CLI before Release 6.1 fail if you use the **-unassign** parameter on Release 6.1 or later DS8000 storage systems with allocated extents. If you use these previous versions of the DS CLI, manually delete or migrate any volumes that use these ranks before you use the **-unassign** parameter.
3. The **-extpool** and **-unassign** parameters cannot be used together.

-extpool *extentpool_ID*
(Optional) Assigns the rank to an extent pool. Accepts either a fully qualified extent pool ID including storage image ID or a shortened version if the **-dev** parameter is used. The shortened version is a four-digit decimal number with no leading zeroes, prefixed with the letter P.

Note: The **-extpool** and **-unassign** parameters cannot be used together.

-quiet
(Optional) Turns off the rank modification confirmation prompt for this command.

rank_ID ... | -
(Required) Specifies one or more ranks to be modified. Accepts either a fully qualified rank ID, or a rank number if the **-dev** parameter is used. A rank number is a four-digit decimal number with no leading zeroes, prefixed by the letter R.

To specify a range of rank IDs, separate the rank IDs with a hyphen.

You must separate multiple rank IDs or ranges of rank IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the chrank command

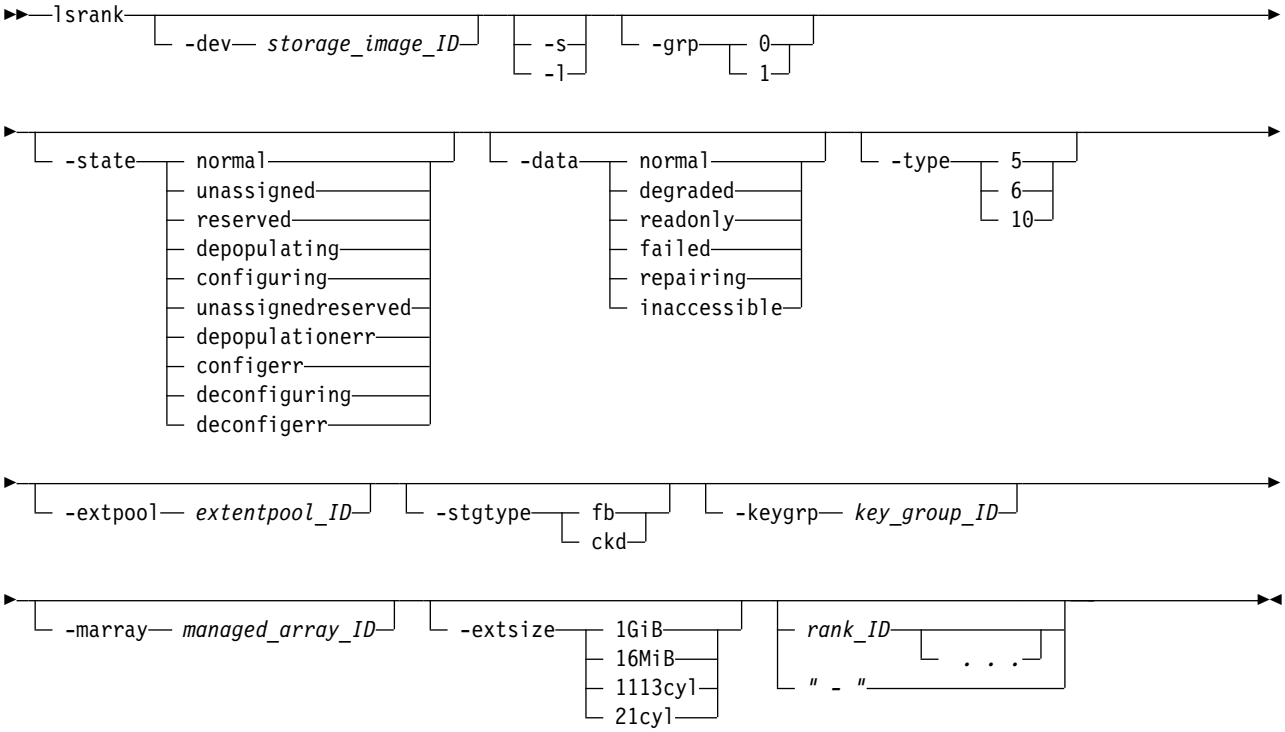
```
dscli> chrank -dev IBM.2107-75FA120 -extpool P101 -quiet R2
```

The resulting output

```
Rank IBM.2107-75FA120/R2 successfully modified.
```

lsrank

The **lsrank** command displays a list of defined ranks in a storage image and status information for each rank.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial

number. The storage image ID is required if you do not specify fully qualified IDs or do not set the *devid* variable in your profile or through the **setenv** command. The storage image ID is also required if HMC is aware of more than one storage image.

Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-s
(Optional) Displays only the rank ID. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output plus the extent pool name, number of extents, extents that are used, and key group for each rank. You cannot use the **-l** and the **-s** parameters together.

-grp 0 | 1
(Optional) Displays only the ranks that belong to the specified rank group. A rank in the unassigned state is not associated with a rank group.

-state normal | unassigned | reserved | depopulating | configuring | unassignedreserved | depopulationerr | configerr | deconfiguring | deconfigerr
(Optional) Displays only ranks in the specified state.

-data normal | degraded | readonly | failed | repairing | inaccessible
(Optional) Displays only ranks in the specified data state.

-type 5 | 6 | 10
(Optional) Displays only ranks of the specified RAID type.

Note: The **-raidtype 6** parameter can be specified for DS8000 models only.

-extpool *extentpool_ID*
(Optional) Displays only ranks in the specified extent pool. An extent pool ID is a four-digit decimal number with no leading zeros, prefixed with the letter P.

-stgtype fb | ckd
(Optional) Displays only ranks that are configured for the specified storage type.

-keygrp *key_group_ID*
(Optional) Displays the key group that this rank uses.

-marray *managed_array_ID*
(Optional) Displays only those ranks that are assigned to the specified managed array.

-extsize
(Optional) Displays only those ranks that have the specified extent sizes.

1GiB
The rank storage type is FB and there is 1 GiB in an extent on this rank.

16MiB
The rank storage type is FB and there is 16 MiB in an extent on this rank.

1113cyl
The rank storage type is CKD and there are 1113 cylinders in an extent on this rank.

21cyl
The rank storage type is CKD and there are 21 cylinders in an extent on this rank.

rank_ID ... | *–*
(Optional) Displays rank information for specified rank IDs. An ID range is defined by two IDs that are separated by a hyphen.

You must separate multiple rank IDs or ranges of rank IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the lsrank command

```
dscli> lsrank  -dev IBM.2107-75FA120 -l
```

The resulting output (The output displays in two parts to avoid text overflow.)

```
ID Group State datastate Array RAIDtype extpoolID extpoolnam           stgtype exts
==========================================================================================
R0  1    Normal Normal    A1     5        P1      pool_large_ckd_wht   ckd      95
R1  1    Normal Normal    A2     5        P3      pool_small_ckd_wht   ckd    20949
usedexts   keygrp  marray   extsize(cap)
========================================
0          -           MA2        1113cyl
0          -           MA3          21cyl
```

Report field definitions

ID Indicates the unique identifier that is assigned to the rank.

Group
Indicates the rank group that the rank is assigned to. A value of 0, 1, or " - " is displayed. If a rank is unassigned, the value that is displayed is " - ".

State
Indicates the current configuration state of this rank ID. One of the following values is displayed:

Normal
Indicates that rank is assigned to an extent pool ID and none of the other state conditions apply.

Unassigned
Indicates that the rank is not assigned to an extent pool ID.

Reserved
Indicates that the rank extents are not eligible for volume assignment.

Depopulating
Indicates that the extents on a rank are not eligible for allocation and the existing allocations are to be moved to another rank in the extent pool by using dynamic extent relocation.

Configuring
Indicates that the rank is being initially configured.

Unassigned Reserved
Indicates that a rank is not assigned to any extent pools, and it is also reserved. A rank with this state will change to Reserved state after it is assigned to an extent pool or changed to the Unassigned state after it is released.

Depopulation Error
Indicates that depopulating the specified rank failed and efforts to depopulate the rank were stopped. (A failed rank can occur when the state of the rank is not valid.) A rank with this state changes to Reserved state if it is reserved, or change to Normal state if it is released.

Configuration Error
Indicates that a rank configuration process failed to complete successfully. This state reflects an internal error condition and not an error in the user's request to create the rank. To correct this state, you must delete the designated rank configuration.

Deconfiguring
Indicates that the rank is being deleted.

Deconfiguration Error
Indicates that a rank removal process did not complete successfully. This state indicates that there is an internal error condition and it is not an indication that there was a user input error. This configuration state is corrected by reissuing the **rmrank** command.

Datastate

Note: A rank is not considered for new extent allocations if it is not in the normal or degraded data state (even if the configuration state is normal).

Datastate specifies the current state of the data extents that are contained by the designated rank ID. One of the following values is displayed:

Normal
A rank is in the normal data state during one of the following configuration states: unassigned, configuring, or configuration error.

Degraded
A rank is in the degraded data state if one or more arrays in the rank are in the degraded data state and none are in the read only, failed, repairing, or inaccessible data states.

Read only
A rank is in the read only data state if one or more arrays in the rank are in the read only data state and none are in the failed, repairing, or inaccessible data states.

Failed
The rank is in the failed data state if one or more arrays in the rank are in the failed data state.

Repairing
A rank is in the repairing data state if one or more arrays in the rank are in the repairing data state and none are in the failed data state.

Inaccessible
A rank is in the inaccessible data state if one or more arrays in the rank are in the inaccessible data state and none are in the failed or repairing data states.

Array
Indicates the array ID that is assigned to the designated rank.

RAIDtype
Indicates the RAID type of the array that is associated with this rank. The value that is displayed is either 5, 6, or 10.

Note: The RAID type 6 is displayed in DS8000 models only.

ExtpoolID
Indicates the extent pool to which the rank is assigned.

Extpoolnam
Indicates the name that is assigned to the extent pool to which the rank is assigned.

Stgtype
Indicates the storage type of the extent pool to which the rank is assigned. The value that is displayed is either FB (fixed block) or CKD (count key data).

Exts
Indicates the number of extents that are contained in the designated rank.

Usedexts
Indicates the number of extents that are allocated to volumes from the designated rank.

keygrp
Indicates the encryption key group. A dash (-) means that either encryption is supported but not used, or encryption is not supported.

marray
Identifies the managed array that is assigned to this rank. This parameter only displays when you specify the **-l** parameter. The following values can display:

"-" If the system does not support managed arrays, a dash (-) displays.

Unknown
If managed arrays are understood by the system, ranks consisting of more than one array are in the Unknown state.

managed_array_ID
If managed arrays are understood by the system, the specified rank has only one array, with only one array site. The range of managed array ID numbers that display is the same as the range of array site numbers, but prefixed with *MA*.

extsize(cap)
Identifies the extent size for the specified rank. There is a direct relationship with the **-extsize** parameter specified for the rank when it is created and the CKD and FB small extent values.

Note: This parameter only displays when you specify the **-l** parameter. The following values can display.

1GiB
The rank storage type is FB and there is 1GiB in an extent on this rank.

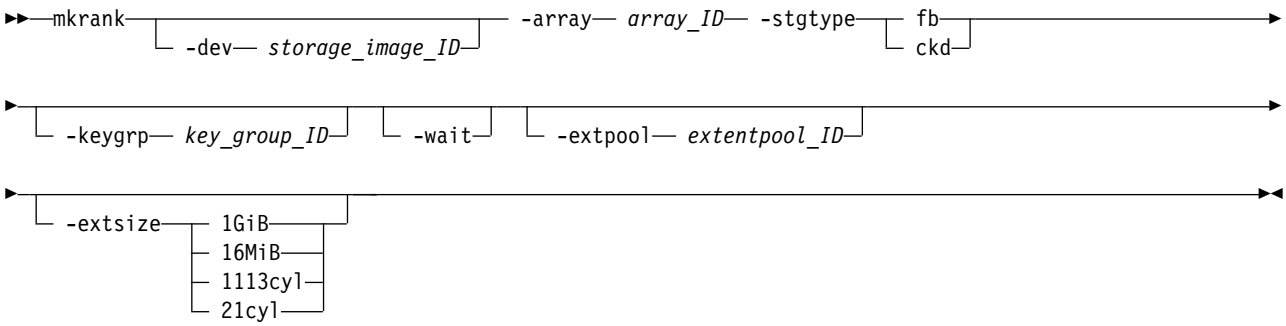
16MiB
The rank storage type is FB and there is 16MiB in an extent on this rank.

1113cyl
The rank storage type is CKD and there are 1113 cylinders in an extent on this rank.

21cyl
The rank storage type is CKD and there are 21 cylinders in an extent on this rank.

mkrank

The **mkrank** command creates one fixed block or count key data (CKD) rank from one array.



Parameters

Notes:

1. The ID is assigned during rank creation based on the current configuration, past configuration changes, and other internal considerations. No algorithm is available to accurately predict the newly created rank ID in every case.

2. Ensure that you specify either the **-wait** or the **-extpool** parameter with the **mkrank** command. Either of these parameters provide notification if the rank configuration fails for any reason.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. For example, IBM.2107-75FA120. The storage image ID is required if you do not specify a fully qualified key group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the storage system is aware of more than one storage image. U

Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-array *array_ID*
(Required) Specifies the ID of the array from which the rank is to be created. An array ID is a four-digit decimal number with no leading zeros, prefixed with the letter A.

-stgtype fb | ckd
(Required) Specifies the type of extent for which the rank is configured for either fixed block or count key data.

-keygrp *key_group_ID*
(Optional) Specifies the encryption key group that this rank uses. The default is zero, which means that no key group is assigned to the rank.

-wait
(Optional) Delays the command response until the rank configuration process completes.

-extpool *extentpool_ID*
(Optional) Specifies the extent pool that contains the created rank extents. If an extent pool is specified, the rank is assigned to the extent pool. Otherwise, the rank state is unassigned. The extent pool you specify must exist and must be compatible with the specified **-stgtype** parameter option. An extent pool ID is a four-digit decimal number with no leading zeros, prefixed with the letter *P*.

Note: If you choose to specify the extent pool ID later, then you must use the **chrank** command.

-extsize
(Optional) Displays only those ranks that have the specified extent sizes.

1GiB
The rank storage type is FB and there is 1 GiB in an extent on this rank.

16MiB
The rank storage type is FB and there is 16 MiB in an extent on this rank.

1113cyl
The rank storage type is CKD and there are 1113 cylinders in an extent on this rank.

21cyl
The rank storage type is CKD and there are 21 cylinders in an extent on this rank.

Examples

Create a CKD small extent rank.

```
dscli> mkrank -array A0 > mkrank -array A3 -stgtype fb -extsize 16MiB
```

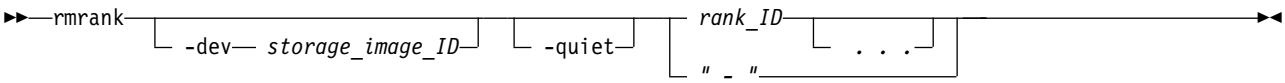
The resulting output

```
Sun Aug 11 02:23:49 PST 2015 IBM DS CLI Device: IBM.2107-75FA120

Rank R4 successfully created
```

rmrank

The **rmrank** command deletes ranks from a storage image. This command is rejected if any volume extents in the rank are being used. In addition, this command formats the drives (DDMs). Until the formatting is done, the associated array cannot be removed.



Parameters

Note: The processing time that is associated with this command can be lengthy and might inhibit your use of the array on which this command is being processed.

When the **rmrank** command is issued, the following processing occurs:

- The rank is unassigned from the array.
- The rank is removed. When this is successful, a message is displayed. This part of the process does not take long; however, the processing that is associated with this command is not complete even though you have received a message that the rank was removed.
- The array is formatted. This processing can take some time. During this process, the array cannot be removed or assigned to another rank. Also, until this process is fully completed, a display of the associated array shows a status of Unavailable. After it is fully completed, the status of the array changes to Unassigned.
- You can check the progress of the **rmrank** command by logging on to another session of DS CLI. Issue the **lsarray** command against the storage image where the rank or ranks are being deleted. If the array state is Unavailable, the **rmrank** command is still processing. When the array state becomes Available, the **rmrank** command is complete.

The following list defines the parameters that are associated with the **rmrank** command:

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for all ranks, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-quiet
(Optional) Turns off the rank removal confirmation prompt for this command.

rank_ID ... | **-**
(Required) Specifies an array of one or more ranks to be deleted. This parameter accepts a fully qualified rank ID, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is a four-digit decimal number with no leading zeroes, prefixed with the letter *R*.

To specify a range of rank IDs, separate the rank IDs with a hyphen.

You must separate multiple rank IDs or ranges of rank IDs with a space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmrank command

```
dscli> rmrank -dev IBM.2107-75FA120 R23
```

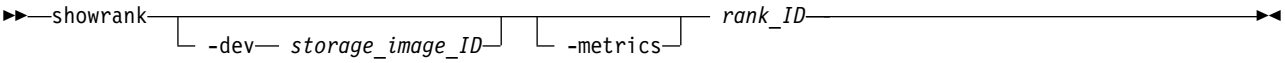
The resulting output

```
Are you sure you want to delete rank R23? [y/n]: Y

Rank R23 successfully deleted.
```

showrank

The **showrank** command displays detailed properties or performance metrics of a rank.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified rank ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-metrics
(Optional) Displays the rank ID and performance statistics for the specified rank.

Note: All performance statistics are an accumulation since the most recent counter wrap or counter reset. Rank performance counters are reset on a power up sequence or by a server failover and failback sequence.

rank_ID
(Required) Specifies the properties for the specified rank. This parameter accepts a fully qualified rank ID, which consists of the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is a four-digit decimal number with no leading zeros, prefixed by the letter *R*.

Example 1

Invoking the showrank command

```
dscli> showrank r4
```

The resulting output

```
ID              R4
SN              -
Group           0
State           Normal
datastate       Normal
Array           A1
RAIDtype        5
extpoolID       P4
extpoolnam      POOL_1GiB_FB
volumes         -
stgtype         fb
exts            352
usedexts        0
widearrays      0
nararrays       1
```

```
trksize        128
strpsize       512
strpesize      0
extsize        16384
keygrp         -
migrating(in)  0
migrating(out) 0
marray         MA2
extsize(cap)   1GiB
```

Report field definitions (-metrics parameter not specified)

ID Indicates the unique ID that is assigned by the system to the rank. The ID includes the storage image ID and the rank ID.

SN Indicates the unique serial number that is assigned to the designated rank ID.

Group Indicates the rank group that the rank is assigned to. One of the following values is displayed: 0, 1, " - ".

Note: " - " is displayed if the rank has not been assigned to an extent pool.

State Indicates the configuration state that is associated with the rank at the time that the report is generated. The following values can be displayed for the rank

Normal
Indicates that a rank is assigned to an extent pool ID and none of the other state conditions apply.

Unassigned
Indicates that a rank is not assigned to an extent pool ID.

Reserved
Indicates that rank extents are not eligible for volume allocation.

Depopulating
Indicates that the extents on a rank are not eligible for allocation and the existing allocations are to be moved to another rank in the extent pool using dynamic extent relocation.

Configuring
Indicates that a rank is in the process of being initially configured. This state indicates that the associated rank transaction has not completed.

Unassigned Reserved
Indicates that a rank is not assigned to any extent pools, and it is also reserved. A rank with this state will change to Reserved state after it is assigned to an extent pool or changed to the Unassigned state after it is released.

Depopulation Error
Indicates that depopulating the specified rank failed and efforts to depopulate the rank were stopped. (A failed rank can occur when the state of the rank is not valid.) A rank with this state will change to Reserved state if it is reserved, or change to Normal state if it is released.

Configuration Error
Indicates that a rank configuration process did not complete successfully. This state indicates that there is an internal error condition and it is not an indication that there was a user input error.

Deconfiguring
Indicates that a rank is in the process of being deleted.

Deconfiguration Error
Indicates that a rank removal process did not complete successfully. This state indicates that there is an internal error condition and it is not an indication that there was a user input error. This configuration state is corrected by reissuing the **rmrank** command.

Datastate
Indicates the current state of the data extents that are contained by this rank ID. The following values can be displayed for the rank.

Normal
Indicates that none of the other data states apply.

Degraded
Indicates that one or more arrays in the rank are in the degraded state.

Read Only
Indicates that one or more arrays in the rank are in the Read Only state.

Failed Indicates that one or more arrays in the rank are in the Failed state.

Repairing
Indicates that one or more arrays in the rank are in the Repairing state.

Inaccessible
Indicates that one or more arrays in the rank are in the Inaccessible state.

Array Indicates the array ID that is assigned to the designated rank.

RAIDtype
Indicates the RAID type (5, 6, or 10) of the array that is associated with the designated rank.

Note: The RAID type 6 is displayed in DS8000 models only.

ExtpoolID
Indicates the extent pool to which the designated rank is assigned.

Extpoolnam
Indicates the extent pool to which the designated rank is assigned.

Volumes
Indicates the volume IDs that have an extent pool value that is allocated on the designated rank.

Stgtype
Indicates the storage type of the extent pool the designated rank is assigned to. Valid values are fixed block and count key data (CKD).

Exts Indicates the number of extents that are contained in the designated rank.

Usedexts
Indicates the number of extents that are allocated to volumes from the designated rank.

Widearrays
Indicates the number of wide arrays that are contained by the designated rank. 0 or 1 are valid values.

Nararrays
Indicates the number of narrow arrays that are contained by the designated rank.

Trksize
Indicates the track size.

Notes:

1. The track size is displayed as a 1 if it is associated with a CKD storage type.
2. The track size is displayed as 128 if it is associated with a fixed block storage type.

Strpsize
Indicates the number of logical tracks in a strip on the designated rank.

Strpesize
Indicates the number of logical tracks in a stripe on the designated rank.

Extsize
Indicates the number of logical tracks in an extent on the designated rank.

Notes:

1. A CKD 1 GB extent contains 16 695 tracks.
2. A fixed block 1 GB extent contains 16 384 tracks.

keygrp
Indicates the key group number. A dash (-) means that either encryption is supported but not used, or encryption is not supported.

migrating(in)
Indicates the number of extents that are migrating into the rank. In other words, this value is the number of migrating extents in the rank that have specified this rank as the target of the migration.

migrating(out)
Indicates the number of extents that are migrating out of the rank.

marray
Identifies the managed array that the specified rank is assigned to using one of the following values.

"-" If the system does not support managed arrays, a dash (-) displays.

Unknown
If managed arrays are understood by the system, ranks consisting of more than one array are in the Unknown state.

managed_array_ID
If managed arrays are understood by the system, the specified rank has only one array, with only one array site. The range of managed array ID numbers that display is the same as the range of array site numbers, but prefixed with *MA*.

extsize(cap)
Specifies the extent size on the specified rank.

1GiB The rank storage type is FB and there is 1GiB in an extent on the specified rank.

16MiB The rank storage type is FB and there is 16MiB in an extent on the specified rank.

1113cyl
The rank storage type is CKD and there are 1113 cylinders in an extent on the specified rank.

21cyl The rank storage type is CKD and there are 21 cylinders in an extent on the specified rank.

Example 2

Invoking the showrank command to show performance metrics

```
dscli> showrank  -dev IBM.2107-75FA120 -metrics R34
```

The resulting output

```
ID             IBM.2107-75FA120/R34
Date           10/11/15 02:23:47
Byteread       10000
Bytewrit       10000
Reads          10000
Writes         10000
Timeread       10000
Timewrite      10000
Dataencrypted  yes
```

Report field definitions (with the -metrics parameter specified)

ID Indicates the unique ID that is assigned by the system to the rank. The ID includes the storage image ID and the rank ID.

Date Indicates the time stamp for the rank performance counters.

Byteread
Indicates the number of rank bytes that are read in 128 KB increments.

Bytewrit
Indicates the number of rank bytes that are written in 128 KB increments.

Reads Indicates the rank read operations.

Writes Indicates the rank write operations.

Timeread
Indicates the rank read response time in 16 millisecond increments.

Timewrite
Indicates the rank write response time in 16 millisecond increments.

dataencrypted
Indicates whether the data stored on the physical media is encrypted.

Extent pool specific commands

There are DS CLI commands that are specifically used to create, modify, and delete extent pools and to display extent pool information.

The following extent pool specific commands are available:

chextpool
Modifies an extent pool.

lsextpool
Generates a report that displays a list of the extent pools in a storage system and the status information on each extent pool in the list.

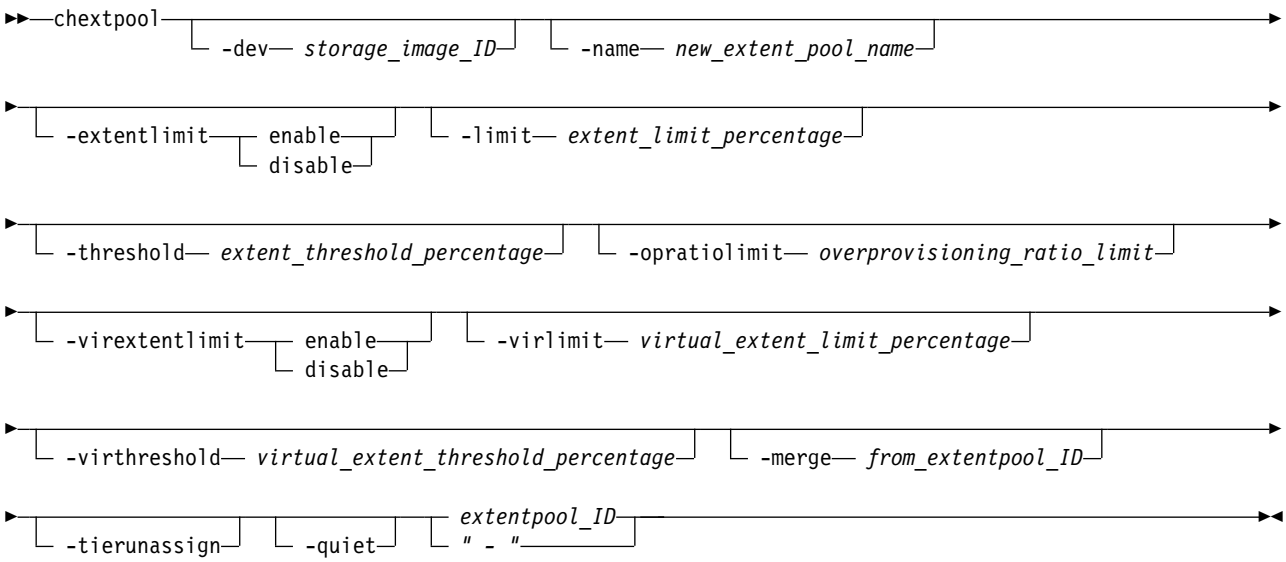
mkextpool
Creates a fixed block or count key data (CKD) storage type extent pool.

rmextpool
Deletes one or more specified extent pools from a storage system.

showextpool
Generates two types of reports. One of the reports displays the detailed properties of a specified extent pool. The other report displays the performance metrics for the specified extent pool.

chextpool

Use the **chextpool** command to modify attributes that are associated with an extent pool, merge two extent pools, or unassign all tier-assigned volumes or extents in the pool.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a value for manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the extent pool, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-name *new_extent_pool_name*
(Optional) Specifies a new name for the extent pool.

-extentlimit enable | disable
(Optional) Specifies that the extent limit function is to be enabled or disabled.

-limit *extent_limit_percentage*
(Optional) Specifies the maximum value of the percentage of allocated real extents that are allowed in this extent pool.

-threshold *extent_threshold_percentage*
(Optional) Specifies threshold as a percentage of the available real extents that is compared to the actual percentage of available real extents. The system issues a warning when this percentage is exceeded. The default threshold is 15. The system triggers a warning when the available capacity is less than 15% of the total pool capacity.

Note: A warning is generated when this percentage is exceeded by capacity allocation to Extent Space Efficient (ESE) volumes.

Note: One third of threshold space is not usable for ESE volume migration.

-opratiolimit *overprovisioning_ratio_limit*
(Optional) Specifies whether overprovisioning is allowed on the extent pool. The default value is 0 (unlimited overprovisioning allowed), whereas 1 specifies that no overprovisioning is allowed (with no assumed production impact resulting from reaching storage capacity limits on the system).

-virextentlimit enable | disable
(Optional) Specifies that the virtual extent limit function is enabled or disabled.

-virlimit *virtual_extent_limit_percentage*
(Optional) Specifies the maximum value of the percentage of allocated virtual extents that are allowed in this extent pool.

-virthreshold *virtual_extent_threshold_percentage*
(Optional) Specifies the minimum threshold percentage of the virtual extents that are available.

-merge *from_extentpool_ID*
(Optional) Specifies an existing extent pool to be merged into the extent pool that is specified by the *extentpool_ID* parameter. When the merge is complete, the ranks that are assigned to *from_extentpool_ID* are reassigned to *extentpool_ID*, and the extent pool that is specified by *from_extentpool_ID* are removed. This parameter cannot be used with any other parameters except the -quiet and -dev parameters.

-tierunassign
(Optional) Specifies that all volume and extent tier assignments in the specified pool are to be canceled. This process does not remove any volumes or extents from the pool. It removes only any specified tier assignments for the volumes and extents in the pool.

-quiet
(Optional) Turns off the confirmation prompt for the **–merge** parameter.

extentpool_ID | –
(Required) Specifies the ID of the extent pool to be changed. This parameter accepts either a fully qualified extent pool ID or a shortened version if the parameter is used, **-dev**. The shortened version is a four-digit decimal number with no leading zeros, prefixed with the letter *P*.

If you use the dash (-), the specified value is read from standard input.

Note: You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the chextpool command

```
dscli> chextpool -name host_4_extpool IBM.2107-75FA120/P21
```

The resulting output

```
Extent Pool IBM.2107-75FA120/P21 successfully modified.
```

Invoking the chextpool command using the virtual extent parameters

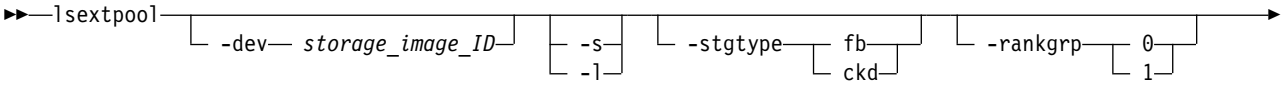
```
dscli> chextpool -dev IBM.2107-75FA120
–virextentlimit enable –virlimit 50 –virthreshold 0 my_extpool
```

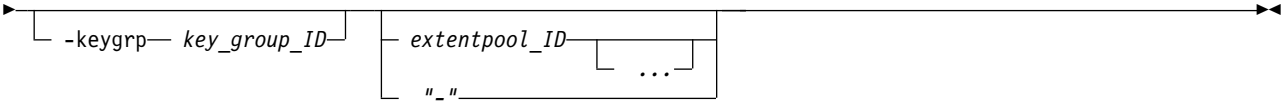
The resulting output

```
Extent Pool IBM.2107-75FA120/P21 successfully modified.
```

lsextpool

The **lsextpool** command displays a list of extent pools in a storage system and status information on each extent pool in the list.





Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. For example, IBM.2107-75FA120. The storage image ID is required if you do not specify a fully qualified key group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the storage system is aware of more than one storage image. U

Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-s
(Optional) Displays only extent pool IDs. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays default output plus the number of ranks and tiers, key group ID, and management status of the extent pools. You cannot use the **-l** and the **-s** parameters together.

-stgtype fb | ckd
(Optional) Displays only extent pools with the specified storage type.

-rankgrp 0 | 1
(Optional) Displays only extent pools in the specified rank group.

-keygrp *key_group_ID*
(Optional) Displays only extent pool of the specified encryption key group.

extentpool_ID ... | -
(Optional) Displays only the extent pools with the specified IDs. An extent pool ID is a four-digit decimal number with no leading zeroes, prefixed by the letter *P*.

To specify a range of extent pool IDs, separate the extent pool IDs with a hyphen.

You must separate multiple extent pool IDs or ranges of extent pool IDs with a space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the lsextpool command

```
dscli> lsextpool -l
```

The resulting output

```
Name               ID  stgtype rankgrp status availstor (2^30B) %allocated available
pool_large_fb_wht  P2    fb      0      below              352        0         352
pool_small_fb_wht  P4    fb      0      below            22584        0       22584

reserved   numvols   numranks     keygrp   numtiers   etmanaged     extsize
 0            0         1           -           1          no          1GiB
 0            0         1           -           1          no         16MiB
```

Report field definitions

Name
Indicates the name you assigned to the extent pool.

ID*
Indicates the system assigned unique identifier for the extent pool object.

Stgtype
Indicates the storage type associated with the extent pool. One of the following types is displayed:

- fb
- ckd

Rankgrp
Indicates the rank group in which the designated extent pool is configured.

Status
Indicates the extent status. One of the following values is displayed:

exceeded
Indicates that the percent of real extents available is less than the real extent threshold.

below
Indicates that the percent of real extents available is greater than the real extent threshold.

full
Indicates that the percentExtents Available is 0.

Availstor (2^30B)
Indicates the available storage for the designated extent pool, in gibibytes (GiB).

%allocated
Indicates the percentage of real extents allocated.

Available
Indicates the maximum number of real extents available for allocation in the designated extent pool.

Reserved
Indicates the real extents reserved in the designated extent pool.

Numvols
Indicates the number of logical volumes that have been configured from the designated extent pool.

Numranks+
Indicates the number of ranks that have been configured in the designated extent pool.

Keygrp+
The encryption key group associated with the extent pool. A dash (-) means that either encryption is supported but not used, or encryption is not supported.

Numtiers+
Indicates the number of tiers in this extent pool.

Etmanaged+
Indicates whether or not the extent pool is managed.

extsize
Specifies the extent size in the extent pool.

Note: In prior models of the DS8000 series, a fixed area on each rank was assigned to be used for volume metadata, which reduced the amount of space available for use by volumes. In the DS8900F family, there is no fixed area for volume metadata, and this capacity is added to the space available for use. The metadata is allocated in the storage pool when volumes are created and is referred to as the pool overhead.

The amount of space that can be allocated by volumes is variable and depends on both the number of volumes and the logical capacity of these volumes. If thin provisioning is used, then the metadata

is allocated for the entire volume when the volume is created, and not when extents are used, so over-provisioned environments have more metadata. For more information, see the *DS8900F Introduction and Planning Guide (GC27-8525)*.

1GiB
The extent pool storage type is FB and there is 1GiB in an extent in this extent pool.

16MiB
The extent pool storage type is FB and there is 16GiB in an extent in this extent pool.

1113cyl
The extent pool storage type is CKD and there are 1113 cylinders in an extent in this extent pool.

21cyl
The extent pool storage type is CKD and there are 21 cylinders in an extent in this extent pool.

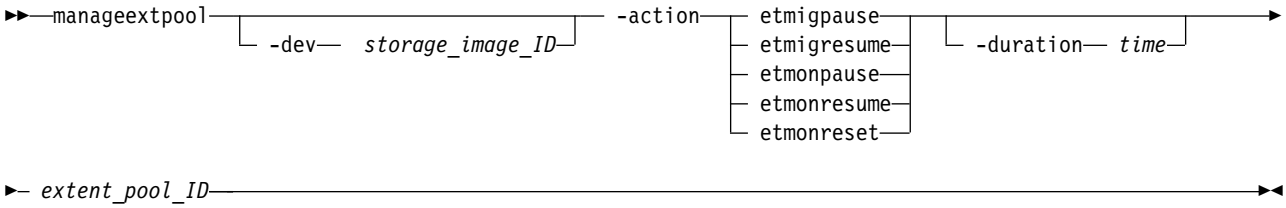
Key:

* Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

manageextpool

The **manageextpool** command starts a process to initiate a change on extent pools.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID or do not set the *devid* variable in your profile. The storage image ID is also required if the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-action etmigpause|etmigresume|etmonpause|etmonresume |etmonreset
(Required) Select one of the following actions:

etmigpause
Specifies that Easy Tier migrations of this storage pool are paused, including migrations that are required to relieve rank bandwidth performance issues. Easy Tier monitoring is unaffected by this action.

etmigresume
Specifies that Easy Tier migrations of this storage pool are resumed. Easy Tier monitoring is unaffected by this action.

etmonpause
Specifies that Easy Tier monitoring of this storage pool is paused. All current Easy Tier migration plans are unaffected, but no new migration plans are formed.

etmonresume
Specifies that Easy Tier monitoring of this storage pool is resumed. Any current Easy Tier migration plans are unaffected.

etmonreset
Specifies that all Easy Tier monitoring data (history), including migration plans, are erased. All new plans are based on new monitoring data.

-duration *time*
(Optional) Specifies the hours of the pause time in ISO 8601 format. For example, **-duration** 24H. The maximum value of the time is a week, which is 168 hours (168H). You can specify this option only with **-action etmigpause** or **etmonpause** parameters.

Note: If you want the duration of the pause to be infinite, you must specify **-duration** 0H. Otherwise, if you do not specify a value with the **-duration** parameter, the default is 168H.

extent_pool_ID
(Required) Specifies the ID of the extent pool that is being managed. Example: IBM.2107-75FA123/P11

Example

Invoking the manageextpool command

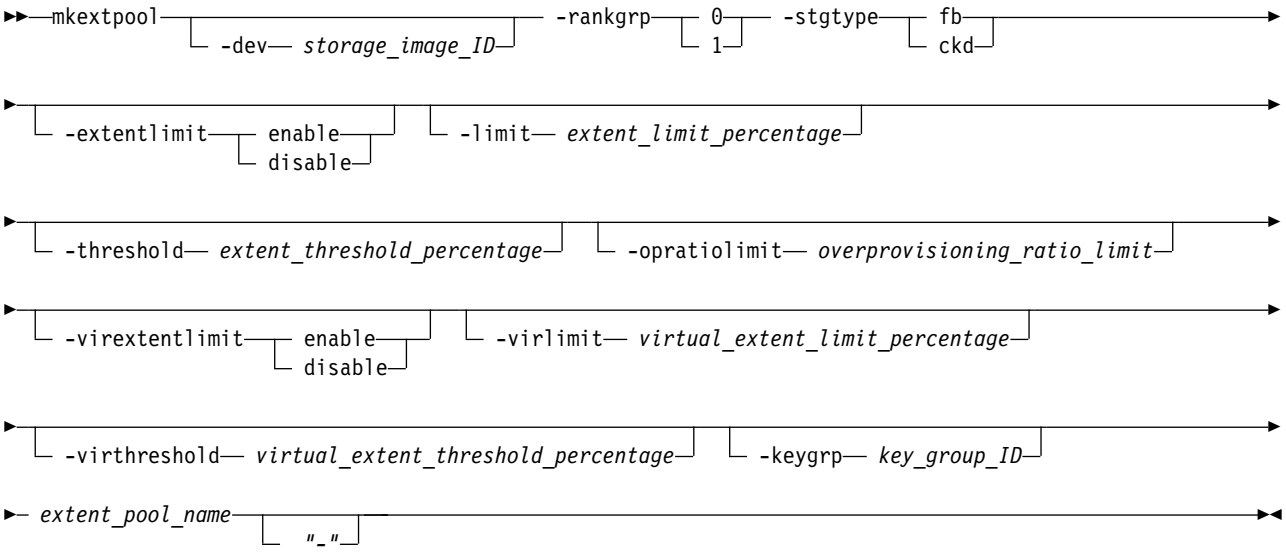
```
dscli  manageextpool –action etmigpause –duration 1H IBM.2107-75FA120/P21
```

The resulting output

```
The etmigpause action for extent Pool IBM.2107-75FA120/P21 has completed.
```

mkextpool

The **mkextpool** command creates a fixed block or count key data (CKD) storage type extent pool.



Parameters

Notes:

1. The storage system assigns the ID during extent pool creation based on the current configuration, past configuration changes, and other internal considerations. No algorithm is available to accurately predict the newly created extent pool ID in every case.
2. An extent pool object is assigned to either rank group 0 or 1. This order allows the extent pool to be managed by storage system server 0 or 1, respectively.
3. Create extent pool objects before creating array and rank objects.

4. Create extent pools of a given type for both rank groups 0 and 1. This method means that volumes that are assigned to a volume group can be spread across both rank groups 0 and 1.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified key group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image.

Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-rankgrp 0 | 1
(Required) Assigns the extent pool to either rank group 0 or 1. Rank group 0 is managed by server 0, and rank group 1 is managed by server 1.

Note: If an extent pool does not exist, you can issue the **chrank** command after an extent pool is created to assign the rank to the extent pool. In addition, you can create extent pools of a given type for both rank groups 0 and 1. This method means that volumes that are assigned to a volume group might be spread across both rank groups 0 and 1.

-stgtype fb | ckd
(Required) Specifies the volume storage type that is contained by this extent pool.

-extentlimit enable | disable
(Optional) Specifies that the extent limit function is enabled or disabled. The default is disable.

-limit *extent_limit_percentage*
(Optional) Specifies the maximum value of the percentage of allocated real extents that are allowed in this extent pool. This value defaults to 100 if not specified.

-threshold *extent_threshold_percentage*
(Optional) Specifies the minimum threshold percentage of the real extents available. When the percentage of the currently available real extents is less than this minimum percentage, notifications are reported. The real extent status is reported as exceeded.

-opratiolimit *overprovisioning_ratio_limit*
(Optional) Specifies whether overprovisioning is allowed on the extent pool. The default value is 0 (unlimited overprovisioning allowed), whereas 1 specifies that no overprovisioning is allowed (with no assumed production impact resulting from reaching storage capacity limits on the system).

-virextentlimit enable | disable
(Optional) Specifies that the virtual extent limit function be enabled or disabled. The default is disable.

-virlimit *virtual_extent_limit_percentage*
(Optional) Specifies the maximum value of the percentage of allocated virtual extents that are allowed in this extent pool.

-virthreshold *virtual_extent_threshold_percentage*
(Optional) Specifies the minimum threshold percentage of the virtual extents available. When the percentage of the currently available virtual extents is less than this minimum percentage, notifications are sent and the virtual extent status is reported as exceeded.

-keygrp *key_group_ID*
(Optional) Specifies the encryption key group that is associated with the specified extent pool. The default is zero, which means that no key group is assigned to the extent pool.

extent_pool_name | **-**
(Required) Specifies your extent pool name, which is limited to 16 characters.

If you use the dash (-), the specified value is read from standard input.

Note: You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the mkextpool command

```
dscli> mkextpool -dev IBM.2107-75FA120 -rankgrp 0 -stgtype fb my_extpool
```

The resulting output

```
Extent pool P2 successfully created.
```

Invoking the mkextpool command using the virtual extent parameters

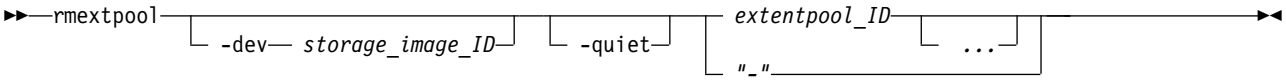
```
dscli> mkextpool -dev IBM.2107-75FA120
-rankgrp 0 -stgtype fb -virextentlimit enable -virlimit 30
-virthreshold 10 my_extpool -keygrp 1
```

The resulting output

```
Extent pool P2 successfully created.
```

rmextpool

The **rmextpool** command deletes extent pools from a storage image.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for all extent pools, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-quiet
(Optional) Turns off the extent pool removal confirmation prompt for this command.

extentpool_ID ... | -
(Required) Specifies the IDs of one or more extent pools to be deleted. A fully qualified extent pool ID is accepted, which consists of the storage image ID, or a shortened version without the storage image ID if the**-dev** parameter is specified. The shortened version is a four-decimal digit number with no leading zeroes, prefixed with the letter *P*.

Note: All rank assignments must be removed before extent pool can be deleted.

To specify a range of extent pool IDs, separate the extent pool IDs with a hyphen.

You must separate multiple extent pool IDs or ranges of extent pool IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmextpool command

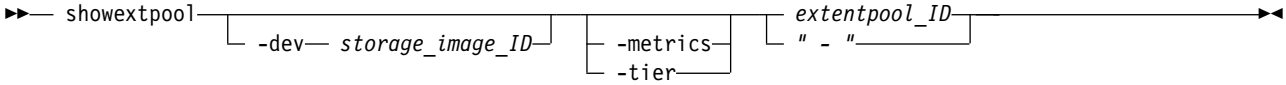
```
dscli> rmextpool IBM.2107-75FA120/P101
```

The resulting output

```
Are you sure you want to delete extent pool IBM.2107.75FA120/P101?
[y/n]: Y
Extent pool IBM.2107-75FA120/P101 successfully deleted.
```

showextpool

The **showextpool** command displays detailed properties, tier information, or performance metrics of an extent pool.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a value for manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified extent pool ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-metrics
(Optional) Displays the extent pool ID and performance metrics for the specified extent pool.

Note: All performance metrics are an accumulation starting from the most recent counter wrap or counter reset. The extent pool performance counters are reset on the following occurrences:

- The storage system is powered-up.
- A server failed and a failover and failback sequence completed.

-tier
(Optional) Displays the extent pool ID and tier information for the specified extent pool.

extentpool_ID | **-**
(Required) Specifies the extent pool that is to be displayed. This parameter accepts a fully qualified extent pool ID, which consists of the storage image ID or an extent pool number without the storage image ID, if the **-dev** parameter is specified. The extent pool number is a four-digit decimal number with no leading zeros, prefixed with the letter *P*. Even-numbered extent pools are associated with rank group 0. Odd-numbered extent pools are associated with rank group 1.

If you use the dash (-), the specified value is read from standard input.

Note: You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 1

Invoking the showextpool command with the -tier option

```
dscli> showextpool -tier p2
Date/Time: March 19, 2017 10:44:07 AM MST IBM DSCLI
Version: 7.7.10.156 DS: IBM.2107-75APK91

Name                    extpool_2
ID                      P2
stgtype                 fb
totlstor (2^30B)        3601
availstor (2^30B)       3048
resvdstor (2^30B)       0
rankgrp                 0
numranks                3
numvols                 11
status                  below
```

```
%allocated                15
%available                84
configured                3601
allowed                   3601
available                 3048
allocated                 553
reserved                  0
configuredCap (MiB/cyl)   442449
allowedCap (MiB/cyl)      442449
availableCap (MiB/cyl)    437829
allocatedCap (MiB/cyl)    4620
reservedCap (MiB/cyl)     0
%limit                    100
%threshold                15
virextstatus              full
%virallocated             100
%viravailable             0
virconfigured             550
virallowed                550
viravailable              0
virallocated              550
virreserved               0
%virextlimit              -
%virextthreshold          -

keygrp                    -
opratio                   1.4
opratiolimit              2
%allocated(ese)           15
%allocated(rep)           0
%allocated(std)           0
%allocated(over)          0
%virallocated(ese)        100
%virallocated(tse)        0
%virallocated(init)       0
%migrating(in)            0
%migrating(out)           0
numtiers                  2
etmanaged                 yes
etmigpauseremain          -
etmonpauseremain          -
etmonitorreset
extsize                   21cyl

2016-07-26T14:00:00+07
=================Tier Distribution================
Tier Cap        (GiB/mod1)  %allocated   %assigned
==================================================
FlashTier2         2121         84           0
ENT                436981       100          14
ENT                436981        100          14
```

Report field definitions

Name
Specifies the name that you assigned to the extent pool.

ID Specifies the system assigned unique identifier for the extent pool object.

stgtype
Specifies the storage type that is associated with the extent pool. One of the following types is displayed:

- fb
- ckd

totlstor (2^30 Bytes)
Specifies the amount of storage that is associated with the extent pool in gibibytes (GiB).

availstor (2^30 Bytes)
Specifies the available storage for the designated extent pool in gibibytes (GiB).

resvdstor (2^30 Bytes)
Specifies the amount of reserved storage for the designated extent pool in gibibytes (GiB).

rankgrp
Specifies the rank group in which the designated extent pool is configured.

numranks
Specifies the number of ranks that are configured in the designated extent pool.

numvols
Specifies the number of logical volumes that are configured from the designated extent pool.

status
Specifies the extent status. One of the following values is displayed:

exceeded
Specifies that the percent of real extents available is less than the real extent threshold.

below
Specifies that the percent of real extents available is greater than the real extent threshold.

full
Specifies that the %Extents available is zero.

%allocated
Specifies the percentage of real extents allocated. A value of 1 - 100 is displayed.

%available
Specifies the percentage of real extents that are available. A value of 1 - 100 is displayed.

configured
Specifies the number of real extents that are contained in the extent pool.

allowed
Specifies the number of real extents that are below the applicable extent limit.

available
Specifies the number of real extents of a specific type that are available for allocation to a logical volume.

allocated
Specifies the number of real extents of a specific type in the extent pool that are allocated to logical volumes or auxiliary volumes.

reserved
Specifies the number of deallocated real extents in the extent pool that are on ranks of the same extent type in the reserved state. In addition, this value is the number of deallocated extents above the applicable extent limit on ranks of the same extent type that are not in the reserved state.

configuredCap (MiB/cyl)
Specifies the total real capacity number. The capacity is reported in MiB for FB extent pool and reports the number of cylinders for CKD extent pool.

allowedCap (MiB/cyl)
Specifies the allowed real capacity number. The capacity is reported in MiB for FB extent pool and reported in number of cylinders for CKD extent pool.

availableCap (MiB/cyl)
Specifies the available real capacity number. The capacity is reported in MiB for FB extent pool and reported in number of cylinders for CKD extent pool.

allocatedCap (MiB/cyl)
Specifies the allocated real capacity number. The capacity is reported in MiB for FB extent pool and reported in number of cylinders for CKD extent pool.

reservedCap (MiB/cyl)
Specifies the reserved capacity. The capacity is reported in MiB for FB extent pool and reported in number of cylinders for CKD extent pool.

%limit
Specifies the maximum percentage of allocated real extents that are allowed in this extent pool.

%threshold
Specifies the minimum threshold percentage of the real extents available. When the percentage of the currently available real extents is less than this minimum percentage, notifications are sent and the extent status is reported as exceeded.

Note: A warning is generated when this percentage is exceeded by capacity allocation to Extent Space Efficient (ESE) volumes.

Note: One third of threshold space is not usable for ESE volume migration.

Virextstatus
Specifies the virtual extent status. One of the following values is displayed:

exceeded
Specifies that the virtual extents available percentage is less than the virtual extent threshold.

below
Specifies that the virtual extents available (**viravailable**) as a percentage of the total extents (**virallowed**) is greater than the virtual extent threshold (**virextthreshold**).

full
Specifies that the available virtual extents is zero.

%virallocated
Specifies the percentage of virtual extents that are allocated compared to the total virtual extents that are allowed. Valid values are 0 - 100.

%viravailable
Specifies the percentage of virtual extents that are available compared to the total virtual extents that are allowed. Valid values are 0 - 100.

Virconfigured
Specifies the number of virtual extents that are configured in the extent pool.

Virallowed
Specifies the number of virtual extents that are below the applicable virtual extent limit.

Viravailable
Specifies the number of virtual extents that are available for allocation to space-efficient volumes.

Virallocated
Specifies the number of virtual extents in the extent pool that are allocated to space-efficient volumes.

Virreserved
Specifies the number of deallocated virtual extents in the extent pool that are on ranks of the same extent type that are in the reserved state. In addition, this value specifies the number of deallocated extents above the applicable extent limit on ranks of the same extent type that are not in the reserved state.

%virextlimit
Specifies the maximum value of the percentage of allocated virtual extents that can be allowed in this extent pool.

Note: If the virtual extent limit is not enabled, a " - " value is displayed.

%virextthreshold
Specifies the minimum threshold percentage of the virtual extents available. When the percentage of the currently available virtual extents is less than this minimum percentage and the virtual extent status is reported as exceeded.

Note: If the virtual extent limit is not enabled, a " - " value is displayed.

keygrp
The key group number that the extent pool is related to. A dash (-) means that either encryption is supported but not used, or encryption is not supported.

opratio
Specifies the overprovisioning ratio, which is the total provisioned capacity divided by total real capacity that is allowable for the volume.

opratiolimit
Specifies whether overprovisioning is allowed on the extent pool. The default value is 0 (unlimited overprovisioning allowed), whereas 1 specifies that no overprovisioning is allowed (with no assumed production impact resulting from reaching storage capacity limits on the system).

%allocated(ese)
Specifies the percentage of allocated real extents on ESE (Extent Space Efficient) volumes compared to the real capacity of the Extent Pool.

%allocated(rep)
Specifies the percentage of allocated real extents on space efficient repository compared to the real capacity of the Extent Pool.

%allocated(std)
Specifies the percentage of allocated real extents on standard volumes compared to the real capacity of the Extent Pool.

%allocated(over)
Specifies the percentage of allocated real extents used for overhead, such as to support space efficient storage, compared to the real capacity of the Extent Pool.

%virallocated(ese)
Specifies the percentage of allocated virtual extents on ESE (Extent Space Efficient) volumes compared to the virtual capacity of the Extent Pool.

%virallocated(tse)
Specifies the percentage of allocated virtual extents on TSE (Track Space Efficient) volumes compared to the virtual capacity of the Extent Pool.

%virallocated(init)
Specifies the percentage of allocated virtual extents being initialized compared to the virtual capacity of the Extent Pool.

%migrating(in)
The percentage of extents migrating into the extent pool compared to the real capacity of the extent pool. (This value is the percentage of migrating extents in the extent pool that have specified this extent pool as the target of the migration.)

%migrating(out)
The percentage of extents migrating out of the extent pool that is compared to the real capacity of the extent pool. (This value is the percentage of migrating extents in the extent pool that have specified this extent pool as the source of the migration.)

numtiers
Specifies the number of tiers in this extent pool.

etmanaged
Specifies whether the extent pool is managed.

etmigpauseremain
Specifies the pause in the Easy Tier migration process. One of the following values is displayed:

0H1M -168H0M
Specifies the time (in hours and minutes) that remain of the pause in the Easy Tier migration process.

infinite
Specifies that Easy Tier migration remains paused until the resume action is submitted.

- (dash)
Specifies that Easy Tier migration is not paused.

unknown
Specifies that the system failed to query the time that remains of the pause.

etmonpauseremain
Specifies the pause in the Easy Tier monitoring process. One of the following values is displayed:

0H1M-168H0M
Specifies the time (in hours and minutes) that remains of the pause in the Easy Tier monitoring process.

infinite
Specifies that Easy Tier monitoring remains paused until a resume action is submitted.

- (dash)
Specifies that Easy Tier monitoring is not paused.

unknown
Specifies that the system failed to query the time that remains of the pause.

etmonitorreset
Easy Tier extent pool monitoring reset date is as follows:

date
Specifies the date of the last Easy Tier monitoring reset in ISO 8601 format: `yyyy-MM-dd'T'HH:mm:ssZ`, where:

yyyy	the year
MM	the month (01-12)
dd	the day (01-31)
T	the single letter T without quotes
HH	the hour (00-23)
mm	the minutes (00-59)
ss	the seconds (00-59)
Z	the time zone offset from UTC [-HHmm | +HHmm]

unknown
Specifies that the date in which Easy Tier monitoring of this extent pool was last reset is not known.

unsupported
Specifies that Easy Tier extent pool management is not supported.

extsize
Specifies the extent size in the extent pool.

Note: In prior models of the DS8000 series, a fixed area on each rank was assigned to be used for volume metadata, which reduced the amount of space available for use by volumes. In the DS8900F family, there is no fixed area for volume metadata, and this capacity is added to the space available for use. The metadata is allocated in the storage pool when volumes are created and is referred to as the pool overhead.

The amount of space that can be allocated by volumes is variable and depends on both the number of volumes and the logical capacity of these volumes. If thin provisioning is used, then the metadata is allocated for the entire volume when the volume is created, and not when extents are used, so over-provisioned environments have more metadata. For more information, see the *DS8900F Introduction and Planning Guide (GC27-8525)*.

1GiB
The extent pool storage type is FB and there is 1GiB in an extent in this extent pool.

16MiB
The extent pool storage type is FB and there is 16GiB in an extent in this extent pool.

1113cyl
The extent pool storage type is CKD and there are 1113 cylinders in an extent in this extent pool.

21cyl
The extent pool storage type is CKD and there are 21 cylinders in an extent in this extent pool.

Report field definitions:

Tier Tier ID

FlashTier0
The tier of high performance flash drives.

FlashTier1
The first tier of high capacity flash drives.

FlashTier2
The second tier of high capacity flash drives.

ENT Enterprise tier; consists of drives with speeds of 10K RPM, 15K RPM, or a combination of drives of both speeds.

NL Nearline tier; consists of high volume SATA or SAS Nearline disk drives.

Unknown
Tier is unknown.

Cap (GiB/mod1)
Capacity of this tier in GiB for FB pools, and mod1 for CKD pools.

%allocated
Capacity of this tier that has been allocated, including assigned capacity, as a percentage of the tier capacity.

%assigned
Capacity that is assigned to this tier as a percentage of the tier capacity. To allow Easy Tier, Easy Tier cooperative caching, and other functions to work correctly, the maximum assigned capacity is limited to 80% of the tier capacity.

Example 2

Invoking the showextpool command to show performance metrics

```
dscli> showextpool -metrics IBM.2107-75FA120/P101
```

The resulting output

ID	Date	real extcap	realext	real allocext	real extconv
IBM.2107-75FA120/P101	10/11/04 02:23:47	10000	10000	10000	10000

dyrelocsource	dyreloctarget
10000	10000

virextcap	virext	dataencrypted
10000	100000	yes

Report field definitions

ID Specifies the system assigned unique identifier for the extent pool object.

Date
Specifies the current time stamp for the extent pool performance counters.

realextcap
Specifies the real extent pool capacity in gibibytes (GiB).

realext
Specifies the number of real extents in the extent pool.

realallocext
Specifies the number of real allocated extents in the extent pool.

realextconv
Specifies real extent conversions.

dyrelocsource
Specifies the number of extents that were sources of a dynamic extent relocation.

dyreloctarget
Specifies the number of extents that were targets of a dynamic extent relocation.

virextcap
Specifies the virtual extent pool capacity in gibibytes (GiB).

virext
Specifies the number of virtual extents in the extent pool.

dataencrypted
Specifies whether the data stored on the physical media is encrypted.

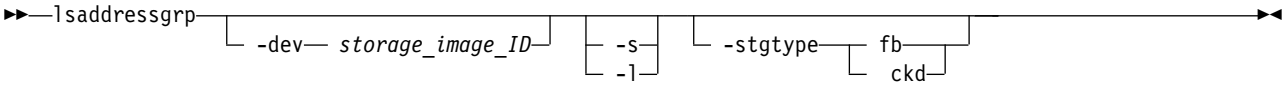
Address group specific commands

The **lsaddressgrp** command displays address group information.

lsaddressgrp
This command generates a report that displays a list of address groups for a storage system. It also provides the status information for each address group in the list.

lsaddressgrp

The **lsaddressgrp** command displays a list of address groups for a storage image and the status information for each address group in the list.



Parameters

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. Displays only the objects for the storage unit specified. The storage image ID is required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-s (Optional). Displays the address group IDs only. You cannot use the **-l** and the **-s** parameters together.

-l (Optional). Displays the default output. You cannot use the **-l** and the **-s** parameters together.

-stgtype fb | ckd
(Optional). Displays only the address groups that are associated with the specified storage type.

Example

The following example represents the headers that are displayed on the output report that is associated with the **lsaddressgrp** command using the **-l** parameter.

Invoking the lsaddressgrp command

```
dscli> lsaddressgrp -dev IBM.2107-75FA120 -l
```

The resulting output

```
ID                 Stgtype Basevolnum Vols LSSs Confgvols
=========================================================
IBM.2107-75FA120/0 fb      0000       4096 16   164096
IBM.2107-75FA120/1 fb      0100       4096 16   164096
IBM.2107-75FA120/2 ckd     0200       4096 16   164096
IBM.2107-75FA120/3 ckd     0300       4096 16   164096
```

Report field descriptions

ID*
Specifies the storage image address group unique identifier. The address number is a single hexadecimal character (0 - 9 or uppercase A - F).

Stgtype
Specifies the type of logical devices that are configured for the specified address group: fb - fixed block and ckd - count key data.

Basevolnum
Specifies the lowest logical volume number in the address group.

Vols
Specifies the number of logical volume numbers that are associated with the address group.

LSSs
Specifies the number of logical subsystems (LSSs) that are configured on the address group.

Confgvols
Specifies the number of logical volumes that are configured on the address group.

Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

Logical control unit specific commands

Commands are referenced for tasks that are associated with IBM Z logical control units.

Use the following logical control unit specific commands:

chlcu Modifies a logical control unit.

lslcu Generates a report that displays a list of logical control units for a storage system and the status information for each logical control unit in the list.

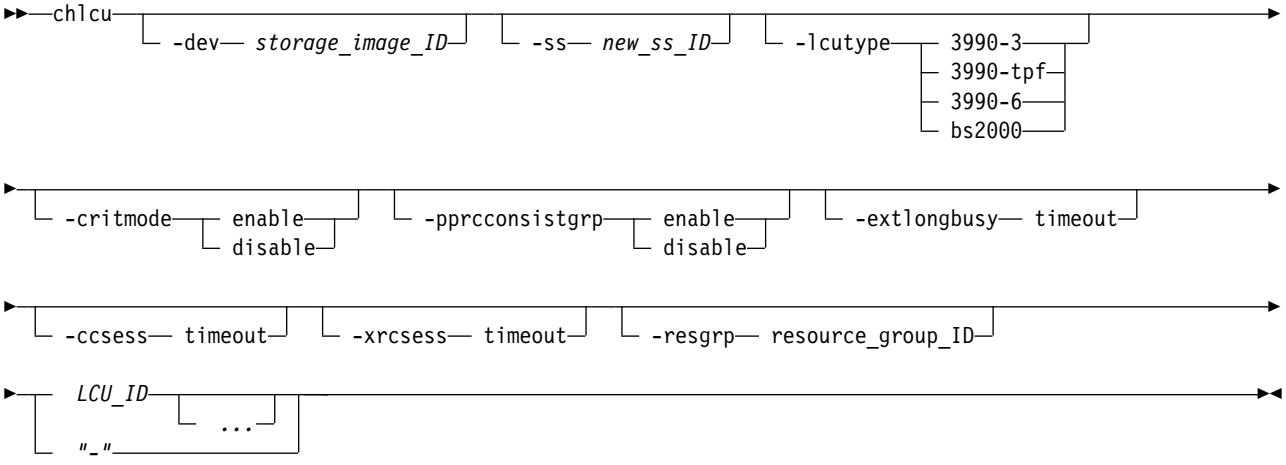
mklcu Creates a logical control unit in a storage system. A logical control unit is configured to represent a grouping of logical CKD volumes.

rmlcu Deletes one or more specified logical control units.

showlcu
Generates a report that displays the detailed properties for the specified logical control unit.

chlcu

The **chlcu** command modifies a logical control unit.



Parameters

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LCU subsystem ID,

do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-ss *new_ss_ID*
(Optional). Specifies your new LCU subsystem ID value (valid range is hexadecimal 0x0001 - 0xFFFF). If this parameter is specified, multiple LCUs are not allowed. The new SSID that you specify replaces the existing SSID value in the initial target LCU ID.

Example: F010

-lcutype 3990-3 | 3990-tpf | 3990-6 | bs2000
(Optional). Changes the target LCUs to the new LCU type:

3990-3 TYPE_3990_MODEL_3

3990-tpf
TYPE_3990_MODEL_3_for_TPF

3990-6 TYPE_3990_MODEL_6

BS2000
TYPE_BS_2000

-critmode enable | disable
(Restricted). Specifies that the critical heavy mode setting in the target LCUs be enabled or disabled. Critical heavy mode controls the behavior of the remote mirror and copy (formerly PPRC) pairs that have a primary logical volume on this LCU and are in an LCU consistency group. See the **mkpprc** command for additional information.

You must have administrator privileges to specify this option.

-pprcconsistgrp enable | disable
(Optional). Specifies that the remote mirror and copy consistency group state setting be enabled or disabled. Any volume that becomes suspended that is associated with the subsystem LSS passes into an extended Long Busy state unless a created consistency group has been received. Otherwise, the volume does not become long busy.

Note: A consistency group volume pair that is suspended will not go into an extended long busy state unless the primary volume is in a full duplex state before the suspension.

-extlongbusy timeout
(Optional). Specifies the time in seconds that an LCU consistency group volume stays long busy after reporting an error that causes a remote mirror and copy (formerly PPRC) suspension if a consistency group has not been created.

-ccsess timeout
(Optional). Specifies the concurrent copy session timeout in seconds setting. This value indicates how long (in seconds) any LCU volume in a concurrent copy session stays long busy before the concurrent copy session is suspended.

Example: 500

-xrcsess timeout
(Optional). Specifies the XRC session timeout value in seconds. This value indicates the time in seconds that any LCU volume in an XRC session stays long busy before the XRC session is suspended. The valid timeout range is 1 - 9999 seconds.

Example: 500

-resgrp *resource_group_ID*
(Optional) Specifies the resource group that the LCUs are assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

LCU_ID ... | -
(Required). Specifies one or more LCUs that are to be modified by this command. A LCU ID is two hexadecimal characters 00 - FE.

You must separate multiple IDs and multiple ID ranges with a space. This parameter accepts a fully qualified LCU ID (which includes the manufacture.machinetype-serial number/ID) or a shortened version if the -dev parameter is specified.

To specify a range of LCU IDs, separate the IDs with a hyphen (-).

If you have specified a new subsystem ID value with the **-ss** parameter, only one LCU ID can be specified.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example: 00-03 08

Example

Invoking the chlcu command

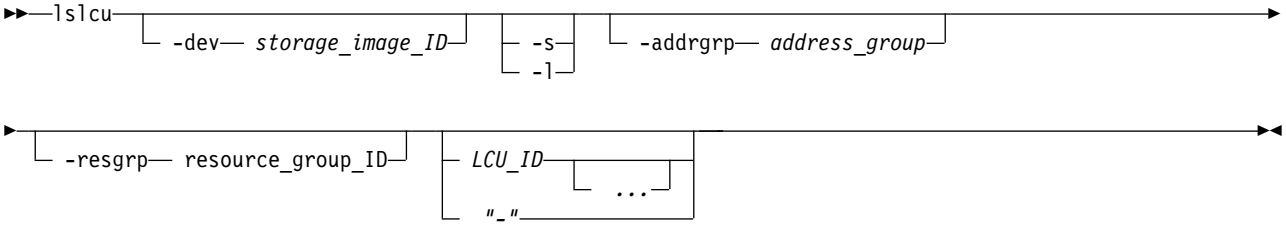
```
dscli> chlcu -dev IBM.2107-75FA120 -critmode enable 00-0F
```

The resulting output

```
LCU 00 successfully modified.
LCU 01 successfully modified.
...
LCU 0F successfully modified.
```

lslcu

The **lslcu** command displays a list of logical control units (LCUs) for a storage image and status information for each logical control unit in the list.



Parameters

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. Displays only the objects for the storage system that is specified. The storage image ID is required for the following conditions:

- A fully qualified LCU ID is not specified.
- The *devid* variable is not set in your profile or specified through the **setenv** command.
- The HMC is aware of more than one storage image.

Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Example IBM.2107-75FA120

-s
(Optional). Displays LCU IDs only. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional). Use this parameter to display the default output plus resource groups. You cannot use the **-l** and the **-s** parameters together.

-addrgrp *address_group*
(Optional). Specifies an address group. Only the LCUs that belong to the specified address group are displayed. An address group is a single character in the range of 0 - 9 or A - F.

-resgrp *resource_group_ID*
(Optional) Displays only the LCUs that are assigned to the specified resource group ID. The resource group ID begins with the letters *RG* and ends with a decimal number.

LCU_ID ... | -
(Optional). Specifies the ID associated with an LCU. An LCU ID is two hexadecimal characters 00 - FE.

To specify a range of LCU IDs, separate the LCU IDs with a hyphen (-).

You must separate multiple LCU IDs or ranges of LCU IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 00-03 08

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lslcu** command by using the **-l** parameter.

Invoking the lslcu command

```
dscli> lslcu -dev IBM.2107-75FA120 -l
```

The resulting output

ID	Group	addr-grp	confg-vols	subsys	con-base-type	resgrp
IBM.2107-75FA120/10	0	0	256	8010	3990-6	RG0
IBM.2107-75FA120/11	1	0	256	8011	3990-6	RG0
IBM.2107-75FA120/12	0	0	256	8012	3990-6	RG0
IBM.2107-75FA120/13	1	0	256	8013	3990-6	RG0

Report field definitions

ID*
Indicates the LCU ID.

The LCU ID is a unique identifier that is assigned to this logical control unit object. The ID includes the storage image ID and the two-digit LCU ID.

Group
Indicates the server that is managing the logical control unit group. The server is identified as either 0 or 1.

Addrgrp
Indicates the address group object of which the logical control unit is a member.

Confgvols
Indicates the number of volumes, or logical devices assigned to the LCU ID. This number includes base CKD volumes and alias CKD volumes.

Subsys
Indicates the value that you assigned, or the default SSID value.

Conbasetype
Indicates the LCU type. The allowable values include the following LCU types:

- 3390-3
- 3390-tpf
- 3390-6 (this value is the default if no value is assigned)
- bs2000

resgrp+
Indicates the resource group ID that the LCU is assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

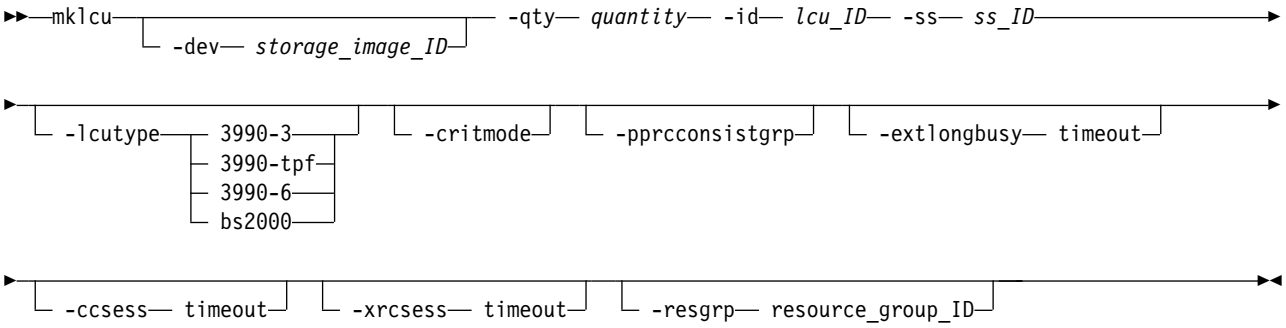
Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

mklcu

The **mklcu** command creates a logical control unit (LCU) in a storage image.



Parameters

Notes:

1. A logical control unit is configured to represent a grouping of logical CKD volumes.
2. Multiple sequential LCU IDs can be created with a single request, but all logical control units must be of the same type and specify the same options.
3. The DS8000 has a 64 KB 256 volume address space that is partitioned into 255 logical subsystem (LSS) units, where each LSS contains 256 logical volume numbers. The 255 LSS units are assigned to one of 16 address groups, where each address group contains 16 LSSs, or 4 KB volume addresses.
4. LCUs are typically created in groups of 16, beginning at LSS address X'x0'.

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required for the following conditions:

- A fully qualified LCU ID is not specified.

- The *devid* variable is not set in your profile or specified through the **setenv** command.
- The HMC is aware of more than one storage image.

Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Example IBM.2107-75FA120

-qty *quantity*
(Required). Specifies the number of LCU IDs to be created. You can specify a quantity value 1 - 255.

This command is rejected if any of the LCU IDs, which are based on the initial LCU ID and the quantity, are currently defined or are outside the range of supported LCU IDs. The valid LCU ID range is 00 - FE.

Example 16

-id *lcu_ID*
(Required). Specifies the LCU ID to be created, or the first LCU ID in a sequence of LCU IDs to be created. An LCU ID is two hexadecimal characters 00 - FE.

Example 00

-ss *ss_ID*
(Required). Specifies the subsystem ID that you assigne. A subsystem ID is four hexadecimal characters 0x0001 - 0xFFFF. If multiple LCU IDs are being created, then the SSID value increments for each additional LCU ID that is created.

If 16 LCUs are created, starting with SSID 0x10, then the SSID values are 0x0010 - 0x001F.

Example 0010

-lcutype 3990-3 | 3990-tpf | 3990-6 | bs2000
(Optional). Specifies that one of the following types of LCU be created:

3990-3 type 3990 model 3

3990-tpf
type 3990 model 3 for tpf

3990-6 type 3990 model 6

bs2000
type bs 2000

-critmode
(Restricted). Specifies that critical heavy mode to enable. Critical Heavy mode controls the behavior of the remote copy and mirror pairs that have a primary logical volume on this LCU. The default value is disable.

You must have administrator privileges to specify this option. See the **mkpprc** command for more notes about the use of this parameter.

Note: If an attempt is made to create the LCU and enable the critical heavy mode but the user does not have the authority to enable the **-critmode** parameter, two messages are displayed when the command is processed:

- One message states that the LCU was successfully created.
- A second message states "Your user ID does not have the authority to perform this operation".

So, the LCU is created but the critical heavy mode is not enabled.

-pprcconsistgrp
(Optional). Specifies a remote mirror and copy consistency group state setting. Any volume that becomes suspended that is associated with the subsystem LSS passes into an extended Long Busy state unless the consistency group that was created previously was received. Otherwise, the volume does not become long busy.

Note: A consistency group volume pair that is suspended will not go into an extended long busy state unless the primary volume is in a full duplex state before the suspension.

-extlongbusy timeout
(Optional). Specifies the time in seconds that an LCU consistency group volume stays long busy after an error is reported that caused a remote mirror and copy suspension if a consistency group has not been created. The default value is 120 seconds.

-ccsess timeout
(Optional). Specifies the **concurrent copy session timeout** parameter as the time in seconds that any LCU volume in a concurrent copy session stays long busy before a concurrent copy session suspends. The valid timeout range is 1 - 9999 seconds. The default value is 300 seconds.

Example 500

-xrcsess timeout
(Optional). Specifies the XRC session timeout parameter as the time in seconds that any LCU volume in an XRC session stays long busy before the XRC session suspends. The valid timeout range is 1 - 9999 seconds. The default value is 300 seconds.

Example 500

-resgrp *resource_group_ID*
(Optional) Specifies the resource group that the LCUs are assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number. The default is *RG0*.

Example

Invoking the mklcu command

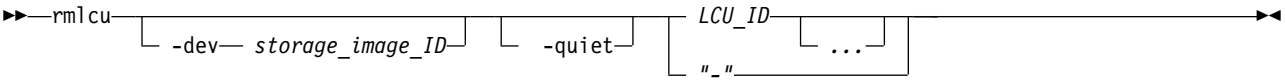
```
dscli> mklcu -dev IBM.2107-75FA120 -qty 16 -id 80 -ss 2300
```

The resulting output

```
LCU 80 successfully created.
LCU 81 successfully created.
...
LCU 8F successfully created.
```

rmlcu

The **rmlcu** command deletes existing logical control units.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for all logical control units, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-quiet
(Optional) Turns off the LCU removal confirmation prompt for this command.

LCU_ID ... | -
(Required) An array of one or more LCUs to be removed. This parameter accepts a fully qualified LCU ID or a shortened version, if the **-dev** parameter is specified. A LCU ID is two hexadecimal characters in the range 00 - FE.

To specify a range of LCU IDs, separate the LCU IDs with a hyphen (-).

You must separate multiple LCU IDs or ranges of LCU IDs with a blank space between each ID or range of IDs.

Example: 00-03 08

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input.

Note: You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmlcu command

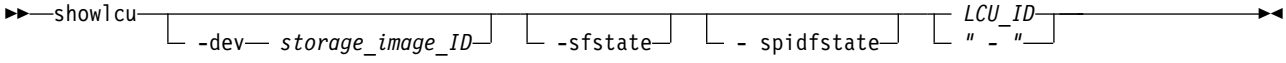
```
dscli> rmlcu -dev IBM.2107-75FA120 00-0F
```

The resulting output

```
Are you sure you want to delete LCU 00-0F ? y/n Y
LCU 00 successfully deleted.
Are you sure you want to delete LCU 01 ? y/n Y
LCU 01 successfully deleted.
...
Are you sure you want to delete LCU 0F ? y/n Y
LCU 0F successfully deleted.
```

showlcu

The **showlcu** command displays the detailed properties of an individual logical control unit (LCU).



Parameters

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the logical control unit, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-sfstate
(Optional) Displays a table at the end of the command output that contains the name, ID, and soft fence state for each volume in the specified LCU.

-spidfstate
(Optional) Displays a table at the end of the command output that contains the name, ID, and SPID (set path group ID) fence state for each volume in the specified LCU.

LCU_ID | -
(Required). Displays the properties for the specified logical control unit. The LCU ID is a 2-digit hexadecimal number in the range of 00 - FE.

Accepts a fully qualified LCU ID, which consists of the storage image ID or a shortened version without the storage image ID, if the **-dev** parameter is specified.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 1

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **showlcu** command.

Invoking the showlcu command

```
dscli> showlcu -dev IBM.2107-75FA120 10
```

The resulting output

ID	Group	Addr-grp	Confg-vols	Subsys	Con-base-type
IBM.2107-75FA120/10	0	1	256	0010	3990-6

Pprc-consist-grp	Xtndlbz-timeout	Ccsess-timout	Xrc-sess-timout	Crit-hvmode	resgrp
Disabled	120 secs	300 secs	300 secs	Disabled	RG0

Report field definitions

ID Indicates the unique identifier that is assigned to this logical control unit object. It includes the storage image ID and the 2-digit LCU ID: 00 - FE.

Group Indicates the server that manages the logical control unit group. The displayed values are 0 or 1.

Addrgrp
Indicates the address group object that the logical control unit is a member of.

Confgvols
Indicates the number of volumes or the logical devices that are assigned to this LCU ID. This number includes base count key data (ckd) volumes and alias ckd volumes.

Subsys
Indicates the subsystem ID that you assigned to this logical control unit. The range of values that you selected from is 0001 - FFFF.

Conbasetype
Indicates the LCU type that you designated for the logical control unit. If you did not designate a type, the default value of 3990-6 is assigned and displayed.

Pprcconsistgrp
Indicates the assigned PPRC consistency group state setting. If you do not designate enabled, the default value of disabled is assigned.

Xtndlbztimout
Indicates the assigned extended long busy timeout value. If you do not designate a value, the default value of 120 seconds is assigned and displayed.

Ccsesstimout
Indicates the assigned concurrent copy session timeout value. If you do not designate a value, the default value of 300 seconds is assigned and displayed.

Xrcsesstimout
Indicates the assigned XRC session timeout value. If you do not designate a value, the default value of 300 seconds is assigned and displayed.

Crithvmode
Indicates whether the critical heavy mode for Remote Mirror and Copy operations is in effect. If you do not designate a value, the default value of Disabled is assigned and displayed.

resgrp Indicates the resource group ID that the LCU is assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

Example 2

If you specify the **-sfstate** parameter, the output includes the soft fence state table.

```
dscli> showlcu -sfstate ef
Date/Time: September 22, 2016 8:43:04 AM MST
IBM.2107-1300861
ID               EF
Group            1
addrgrp          E
confgvols        4
subsys           0x1111
conbasetype      3990-6
pprcconsistgrp   Disabled
xtndlbztimout    120 secs
ccsesstimout     300 secs
xrcsesstimout    300 secs
crithvmode       Disabled
resgrp           RG0
============Soft Fence State============
Name ID   sfstate
=======================
-    EFFC Disabled
-    EFFD Disabled
-    EFFE Disabled
ffff EFFF Disabled
dscli>
```

Report field definitions (-sfstate parameter is specified)

Name The user-assigned nickname for this volume object.

ID The unique identifier that is assigned to this volume object. A volume ID is four hexadecimal characters (0x0000 - 0xFEFF).

sfstate The soft fence state. The following values are possible.

Enabled
The host has set this volume to the soft fence state.

Disabled
The host has not set this volume to the soft fence state.

N/A The host cannot set this volume to the soft fence state. For example, an alias volume.

Report field definitions (-spidfstate parameter is specified)

Name The user-assigned nickname for this volume object.

ID The unique identifier that is assigned to this volume object. A volume ID is four hexadecimal characters (0x0000 - 0xFEFF).

spidfstate
The soft fence state. The following values are possible.

Enabled
The volume is set to the SPID fence state by the host.

Disabled
The volume is not set to the SPID fence state by the host.

N/A The volume is not capable of being set into the SPID fence state by the host. For example, an alias volume.

CKD logical volume specific commands

Commands are referenced for tasks that are associated with IBM Z count key data (CKD) logical volumes.

The following CKD logical volume specific commands are available:

chckdvol
Modifies the nickname that you assigned to the count key data (CKD) base volume.

initckdvol
Releases extents from a space-efficient logical volume, and removes (erases) data from standard volumes.

lsckdvol
Generates a report that displays a list of count key data (CKD) base and alias volumes in a storage system and the status information for each volume in the list.

manageckdvol
Initiates a change on count key data (CKD) volumes by executing a process.

mkaliasvol
Creates IBM Z CKD alias volumes (referred to as parallel access volumes or PAVs) in a storage system.

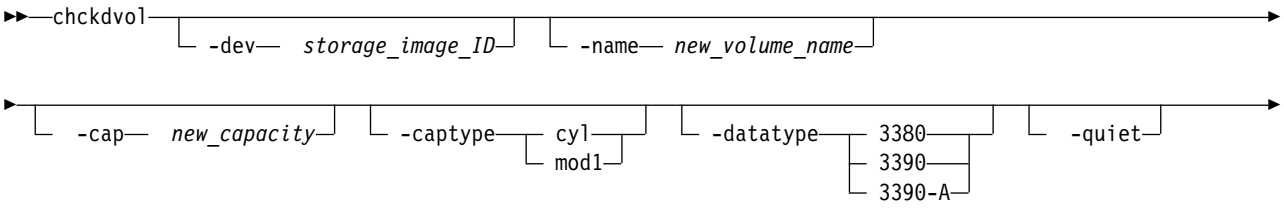
mkckdvol
Creates IBM Z count key data (CKD) base CKD volumes in a storage image.

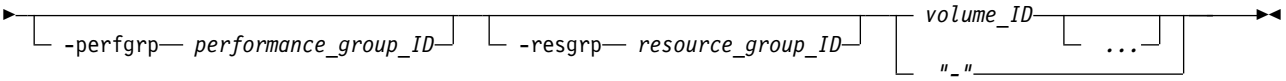
rmckdvol
Deletes one or more specified count key data (CKD) base or alias volumes from a storage system.

showckdvol
Generates two reports. One report displays the detailed properties of a specified count key data volume. The other report displays the performance metrics for specified count key data volume.

chckdvol

The **chckdvol** command modifies the attributes that are associated with a count key data (CKD) base volume.





Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-name *new_volume_name*
(Optional) User specified nickname for this CKD base volume. This nickname should not exceed 16 characters. It can contain one of the following wildcards:

- (#d) - insert volume_ID (decimal format)
- (#h) - insert volume_ID (hexadecimal format)

-datatype 3380|3390|3390-A
(Optional) Specifies the volume data type.

For Extended Addressing Volumes (EAV), which have a capacity greater than 65520 cylinders, the data type 3390-A must be specified. You can also specify the data type 3390-A for volumes smaller than 65520 cylinders.

Notes:

1. You can specify the data type 3390–A for DS8000 models only.
2. You must ensure that the volume data type is compatible with the host systems that access this volume.
3. You cannot use the **-datatype** parameter and the **-cap** parameter in the same command.

-cap *new_capacity*
(Optional) Specifies the quantity of CKD cylinders that you want to allocate to the specified volume. You can also use the **-captype** parameter to change the unit type of the capacity to mod1 units.

- 3380 volumes cannot be expanded.
- For 3390 volumes, the *new_capacity* value can be specified in increments of 1 in the range 1 - 65520 (849KiB to 55.68 GiB).
- For 3390–A volumes, for values less than 65520 cylinders, *new_capacity* is specified in increments of 1. Capacities greater than 65520 cylinders are specified in increments of 1113. The maximum volume size varies and depends on DS8000 model and type.

Notes:

1. Check your operating system documentation to ensure that volume expansion is supported before proceeding with the expansion.
2. If you expand a 3390 volume to an Extended Addressing Volume (EAV), the data type is changed to 3390-A.
3. Attempting to reduce the size returns an error and causes the transaction to fail.
4. You cannot use the **-datatype** parameter and the **-cap** parameter in the same command.

-captype cyl | mod1
(Optional) Specifies the unit type of the capacity given by using the **-cap** parameter. The default is cyl. A mod1 unit is equivalent to 1113 cylinders, and 1263.28 cylinders is equivalent to 1 GiB. You must specify the **-cap** parameter to specify the **-captype**.

-quiet
(Optional) Turns off the CKD volume modification confirmation prompt for this command.

-perfgrp *performance_group_ID*
(Optional) Specifies the performance group ID that the volumes are assigned to. The performance group ID begins with the letters *PG*, and valid performance groups numbers are 0, and 16-31. The default is *PG0*.

-resgrp *resource_group_ID*
(Optional) Specifies the resource group that the volumes are assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

volume_ID ... | -
(Required) An array of one or more CKD base volume IDs or volume ID ranges to modify.

A volume ID range is defined by two volume IDs that are separated by a dash. Multiple volume IDs or volume ID ranges must be separated with a blank space between each ID.

Example: 0100-010F 0180-018F 0120

The volume ID format is four hexadecimal characters *LLVV* that represent the following values:

LL Specifies the logical control unit number, 00 - FE

VV Specifies the volume number, 00 - FF

You must fully qualify the volume ID with manufacturer, machine type, and serial number if you do not use the **-dev** parameter.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) if you are in the DS CLI interactive mode.

Example

Invoking the chckdvol command

```
dscli> chckdvol -dev IBM.2107-75FA120  -perfgrp PG18 0100
```

The resulting output

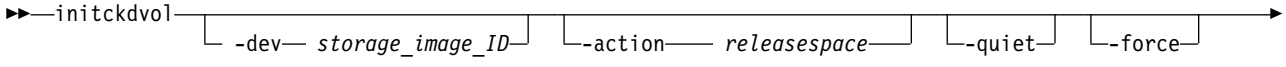
```
CKD volume 0100 successfully modified.
```

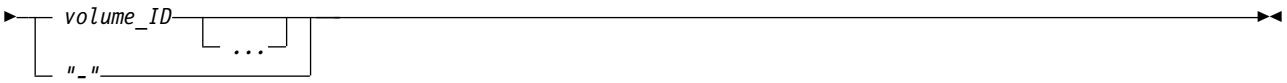
initckdvol

The **initckdvol** command initializes a logical volume and releases extents from a space-efficient logical volume.

For space efficient logical volumes, this command is used to release space. For TSE volumes, it releases tracks in the repository, reducing the repository allocated space. For ESE volumes, it releases extents in the extent pool being used, reducing the allocated extents.

For example, if a space-efficient logical volume has data that is stored on it that is no longer needed, use the **initfbvol** command to free the extents/tracks that were assigned to this logical volume. This allows the extents/tracks to be reused by other space-efficient logical volumes.





For example, if a space-efficient logical volume is used as a FlashCopy target volume and the data that is stored on these tracks are no longer needed, use the **initckdvol** command to free the extents that were assigned to this logical volume. This allows the extents to be reused by other space-efficient logical volumes.

Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a value for manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-action *releasespace*
(Optional) Specifies that you want the system to release the repository space held by the designated space-efficient volume back to the repository. (The repository is the physical extents that provision the virtual extents for virtual space volumes.)

-quiet
(Optional) Turns off the confirmation prompt for the **-action** parameter.

-force
(Optional) Forces the action.

Note: For release 8.3 or later, if **-force** is not specified, the request fails when the volume is in a path group or in any kind of Copy Services relationship.

volume_ID ... | -
(Required) The volume ID format is four hexadecimal characters *LLVV*, that represent the following values:

LL Specifies the logical control unit number, 00 - FE.

VV Specifies the volume number, 00 - FF.

You must fully qualify the volume ID with manufacturer, machine type, and serial number if you do not use the **-dev** parameter.

To specify a range of volume IDs, separate the volume IDs with a dash.

You must separate multiple volume IDs or ranges of volume IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive command mode.

Example

Invoking the initckdvol command

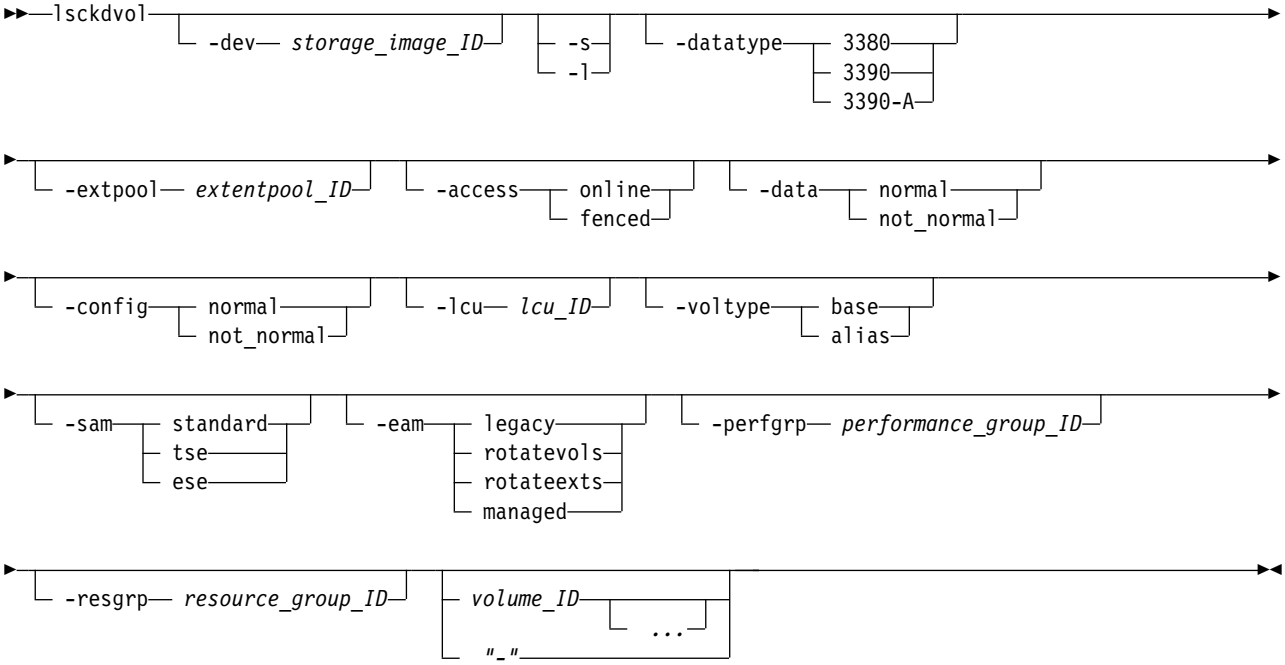
```
dscli> initckdvol -dev IBM.2107-75FA120 -action releasespace 0101
```

The resulting output

```
CMUC00338W initckdvol: Are you sure that you want to submit the command
releasespace for the CKD volume 0101?[Y/N]:y
CMUC00341I initckdvol:: 0101: The command releasespace has
completed successfully.
```

lsckdvol

The **lsckdvol** command displays a list of count key data (CKD) base and alias volumes in a storage image and status information, including Safeguarded Copy information for each volume in the list.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-s
(Optional) Displays only volume IDs associated with the storage image ID. You cannot use the **-s** and the **-l** parameters together.

-l
(Optional) Displays the default output plus additional attributes that are identified as long output in the Report field definitions list. You cannot use the **-s** and the **-l** parameters together.

-datatype 3380 | 3390 | 3390-A
(Optional) Specifies that you want the system to display only volumes that meet the designated volume data type.

Note: The data type 3390–A can be displayed for DS8000 models only.

-extpool *extentpool_ID*
(Optional) Specifies that you want the system to display only volumes that are associated with the designated extent pool.

-access online | fenced
(Optional) Specifies that you want the system to display only volumes with the specified access state.

-data normal | not_normal
(Optional) Specifies that you want the system to display only the volumes that meet the criteria of the designated data state.

-config normal | not_normal
(Optional) Specifies that you want the system to display only the volumes that meet the criteria of the designated configuration state.

-lcu *lcu_ID*
(Optional) Specifies that you want the system to display only volumes with IDs that contain the specified logical control unit ID. Each logical control unit can contain up to 256 volumes. The LCU ID is a 2-digit hexadecimal number in the range of 00 - FE for the DS8000.

Note: Including the **-lcu** parameter can significantly reduce response time because the entire storage system does not have to be queried.

-voltype base | alias
Specifies the type of CKD volumes that you want the system to display.

-sam standard | tse | ese
(Optional) Specifies that you want the system to display only volumes that meet the criteria of the storage allocation method as follows:

standard
Volumes that are fully allocated with real extents.

tse Track space-efficient logical volumes that contain a set of virtual extents that are associated with the space-efficient storage in the same extent pool.

Note: Track space-efficient (TSE) repositories are required to create TSE volumes. However, starting with DS8000 Release 8, TSE repositories can no longer be created.

ese Extent space efficient logical volumes that are provisioned with a set of virtual extents that are associated with the space efficient storage in the same extent pool.

-eam legacy | rotatevols | rotateexts | managed
(Optional) Specifies that you want the system to display only volumes that meet the criteria of the extent allocation method as follows:

legacy Specifies that the volume was created before the use of the current algorithm.

rotateexts
Specifies that volumes were created using storage-pool striping.

rotatevols
Specifies that each successive logical volume that is created is allocated on the next available rank in the extent pool.

managed
Specifies that the volume is in an extent pool managed by Easy Tier.

-perfgrp *performance_group_ID*
(Optional) Specifies the performance group ID that the volumes are assigned to. The performance group ID begins with the letters *PG*, and valid performance groups numbers are 0, and 16-31. The default is *PG0*.

-resgrp *resource_group_ID*
(Optional) Displays only the volumes that are assigned to the specified resource group ID. The resource group ID begins with the letters *RG* and ends with a decimal number.

volume_ID ... | -
(Optional) Displays volumes with the specified IDs. The volume ID format is four hexadecimal characters *LLVV* that represent the following values:

LL Specifies the logical control unit number, 00 - FE

VV Specifies the volume number, 00 - FF

To specify a range of volume IDs, separate the volume IDs with a dash (-).

You must separate multiple volume IDs or ranges of volume IDs with a blank space between each ID or range of IDs.

Example: 0100-010F 0180-018F 0120

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI show commands, the results are shown in table format for clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lsckdvol** command using the **-l** parameter.

The following example is based on the output results for a volume with 3340 cylinders.

Invoking the lsckdvol command

```
dscli> lsckdvol -dev IBM.2107-1300861 -l 1410
```

The resulting output

Name	ID	acc state	data state	config state	device MTM	volser	data type
My_ volume _1410	1410	Online	Normal	Normal	3390-9	A03976	3390

voltype	orgbvols	extpool	sam	cap (cyl)	cap (2^30B)	cap (10^9B)
CKD base	-	P2	standard	3340	2.6	2.8

reqcap (cyl)	eam	perfgrp	resgrp
3339	legacy	PG18	RG0

safeguardedcap (cyl)	safeguardedloc	usedsafeguardedcap (cyl)	safeguarded
1234	mypool_1	0	yes
1234	mypool_1	0	yes
1234	mypool_1	0	yes
1234	mypool_1	0	14

Report field definitions

Name
Indicates the nickname that you assigned to the designated volume object.

ID Indicates the unique identifier that is assigned to the designated volume object

Accstate (access state)
One of the following designations can be displayed:

Online
Indicates that the logical volume is accessible to a host.

Fenced
Indicates that the logical volume is in the volume fenced state and is not accessible to the host.

" - "
Indicates that the logical volume is designated as a CKD alias volume.

Datastate
One of the following designations can be displayed:

Normal
Indicates that none of the other data states apply. The access state is Online.

Pinned
Indicates that none of the other data states apply and the logical volume has one or more pinned non-retryable tracks. The access state is Online.

Read only
Indicates that the logical volume is read only because one or more extents on the logical volume are on a rank in the read only data state. The access state is Online.

Inaccessible
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the inaccessible data state. The access state is Fenced.

Indeterminate data loss
Indicates that the following data states do not apply and that one of the following conditions has occurred:

Data states that do not apply:

- Rank failed
- Rank repairing
- Rank repaired
- Global inaccessible
- Global lost data

Conditions - one of the following occurred:

- Committed write data was lost before it was destaged and the track identifiers that are associated with the data are unknown.
- Data was lost that indicated extents on the logical volume were active FlashCopy targets.

The access state is fenced.

Rank failed
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the failed data state. The access state is Fenced.

Rank repairing
Indicates that one or more extents that are associated with the logical volume are on ranks that are in the repairing data state. The access state is Fenced.

Rank repaired
Indicates that one or more extents that are associated with the logical volume are on ranks that were in the repairing state, but are not in the repairing state now. The access state is Fenced.

Global inaccessible
Indicates that the global metadata that is associated with the logical volume configuration is inaccessible. Some of the data that is associated with the logical volume might be inaccurate. The access state is Fenced.

Global lost
Indicates that global metadata that is associated with the logical volume configuration has been lost. As a result, some of the data that is associated with the logical volume might be inaccurate. The access state is fenced.

NVS data inaccessible
Indicates that active NVS data is inaccessible for one or more logical volumes of an LSS group. The logical volumes in the LSS group cannot be made accessible. The access state is Fenced.

" - "
Indicates that the logical volume is designated as a CKD alias.

Configstate
One of the following configuration states are displayed:

Normal
Indicates that there are no logical volume configuration operations in progress, and the volume is not being deconfigured, merged, or migrated.

Configuring
Indicates that the logical volume is in the process of being configured for the first time.

Reconfiguring
Indicates that the logical volume is in the process of allocating or deallocating extents due to a modification of the requested capacity attribute after initial creation.

Deconfiguring
Indicates that the logical volume is in the process of being deleted.

Configuration error
Indicates that the initial configuration did not complete successfully. This state reflects an internal error condition and not an error in the request to create the volume. If you have a volume in this state, use the **rmfbvol** command to delete each volume listed with the configuration state of "configuration error".

Merging
Indicates that the volume is in the process of merging. For example, merging from one extent pool to a different extent pool.

Migrating
Indicates that the volume is in the process of migrating, or waiting to be migrated.

Migration Cancelled
Indicates that the volume was in the process of migrating and then the 'migcancel' action of the **manageckdvol** command was issued, leaving some of the extents waiting to be migrated in the source pool while other extents have already migrated to the target pool. Migration has stopped, and cannot be resumed. If you have a volume in this state, try to migrate it again to the original source or target extent pool.

Migration Paused
Indicates that the volume was in the process of migrating and then the 'migpause' action of the **manageckdvol** command was issued. Migration has stopped, but can be resumed.

Migration Error
Indicates that the volume migration process failed to complete successfully. This state reflects an internal error condition and not an error in the request of the user to migrate a volume. If you have a volume in this state, try to migrate it again to the original source or target extent pool.

Reconfiguration error
Indicates that the reconfiguration request did not complete successfully.

Deconfiguration error
Indicates that a request to delete a volume did not complete successfully. This state reflects an internal error condition and not an error in the request to remove the volume. To correct this state, you must reissue the **rmfbvol** command for the designated volume.

Transposition Error
Indicates that the volume is in an extent pool that was unsuccessfully merged. This state reflects an internal error condition. **Corrective action**: Use the **chextpool** command with the **-merge** parameter again to redrive the merge extent pool and to correct this state.

deviceMTM
One of the following device MTMs is displayed:

- 3380-2
- 3380-3
- 3390-3
- 3390-9
- 3390-A

Device MTM is determined by the CKD base volume data type and provisioned capacity (in cylinders).

Note: The device MTM 3390–A can be displayed for DS8000 models only.

Volser
Indicates the base CKD volume serial number written by the user at track address 0x0000, record 3.

Datatype
Indicates the volume data type setting.

Voltype
Indicates that the logical volume is one of the following types:

- CKD base
- CKD alias

Orgbvols (original base vol)
One of the following original base volumes is specified:

- Indicates the original base CKD volume ID to which this CKD alias volume is assigned.
- " - " is displayed for a CKD base volume type.

Note: For a CKD Alias type volume, the base and alias volume IDs are of a common LCU ID.

Extpool
Indicates the extent pool ID. Volume extents are allocated from this extent pool ID.

Note: Volumes that belong to an encrypted extent pool are encrypted. You can see the key group of an extent pool by using the **lsextpool -l**, or **showextpool** commands.

SAM
Indicates the storage allocation method. The following values are displayed:

standard
Designates that the system fully allocated the volume with real extents at volume creation time.

tse
Designates that a track space-efficient logical volume contains a set of virtual extents that are associated with the space-efficient storage in the same extent pool. Physical space for a given logical track on a track space-efficient logical volume is dynamically allocated and deallocated from the repository in the space-efficient storage.

ese
Designates that an extent space efficient logical volume is provisioned with a set of virtual extents that are associated with the space efficient storage in the same extent pool. Physical space for an extent space efficient logical volume is dynamically allocated and de-allocated from the extent pool.

Cap (cyl)
Indicates the quantity of volume CKD cylinders that are available for host system access.

Cap (2^30B)
Indicates the size of volume that is available for host system access in gibibytes (GiB).

Cap (10^9B)
Indicates the size of volume that is available for host system access in decimal gigabytes (GB).

Reqcap (cyl)
Indicates the requested quantity of volume CKD cylinders (for example, 3339).

EAM
Indicates the extent allocation method that will be used if the volume is migrated or expanded.

legacy
Indicates that the volume was created before the use of the current algorithm.

rotateexts
Indicates that the extents for each new logical volume are allocated across all available ranks, and is also known as storage-pool striping. This value is the default.

rotatevols
Indicates that the extents for each new logical volume are allocated from each successive rank. This means that the extents for a particular volume will be allocated from one rank, while the extents for the next volume will be allocated from the next successive rank, and so on.

managed
Indicates that the extents are currently managed by Easy Tier, and the extents for any new volumes are initially allocated across all available ranks in the lowest tier of storage.

" - "
A dash (-) value is displayed if the extent allocation method does not apply , for example if the volume is a track space efficient (TSE) volume.

perfgrp
Indicates the performance group ID that the volume is assigned to. The performance group ID begins with the letters *PG* and ends with a decimal number.

resgrp
Indicates the resource group ID that the volume is assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

safeguardedcap(cyl)
Specifies the amount of safeguarded backup capacity that is reserved for safeguarded copies.

safeguardedloc
Specifies the pool where the safeguarded copies are written to.

usedsafeguardedcap(cyl)
Specifies the amount of physical pool capacity that is used by the safeguarded copies.

safeguarded
Specifies whether safeguarded copy is enabled for the specified volume.

manageckdvol

The **manageckdvol** command starts a process to initiate a change on count key data (CKD) volumes.

Use this command to define safeguarded backup capacity for CKD volumes for which you want to enable for Safeguarded Copy.

Use Copy Services Manager to define a safeguarded copy schedule to create backups on a regular basis, such as hourly or daily.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-action
migstart|**migcancel**|**migpause**|**migresume**|**sfdisable**|**spidfdisable**|**tierassign**|**tierunassign**|**etmonpause**|**etmonresume**|**etmonreset**|**mksafeguardedcap**|**rmsafeguardedcap**|**migsafeguardedcap**

|migsafeguardedcapcancel|migsafeguardedcappause|migsafeguardedcapresume
(Required) Specifies that one of the following actions is to be performed:

migstart
Initiates volume migration on the specified volumes that are in the "normal" or "canceled" state. Volumes are placed into the migrating state. Volumes that are in the canceled state must have the original source or the destination extent pool as the value of the **-extpool** parameter.

migcancel
Cancels volume migration on the specified volumes that are in the migrating state. Volumes that have not yet started migrating data are defined to a "normal" state. Volumes that have started migrating data are define to a "canceled" state.

migpause
Pauses volume migration on the specified volumes that are in the migrating state. Volumes that have not yet started migration or that are in the middle of migration are put in the "paused" state.

migresume
Resumes volume migration on the specified volumes that are in the "paused" state.

sfdisable
Sends a soft fence reset command to each specified volume. This action cannot be used with any other parameter.

spidfdisable
Sends a set path group ID (SPID fence reset command to each specified volume. This action cannot be used with any other parameters.

tierassign
Initiates the assigning volume action on the specified volumes to the specified tier. The **-tier** option is required with this action.

Note: By assigning a volume to a tier, a process that migrates the data in the volume to the specified tier begins. You can check the progress of the migration by using the **showfbvol** or **showckdvol** command with the **-tier** parameter. However, there is a maximum capacity that can be assigned to each tier in a pool. If this maximum capacity is reached, then all of the volumes that are still migrating to that tier, because of assignment, pause their migrations. You can check how much total capacity has been assigned to a specific tier in a pool by using the **showextpool** command with the **-tier** parameter. If the maximum was reached, you can remedy the problem either by unassigning other volumes from the tier or by adding more capacity to the tier. Either method automatically restarts any paused migrations.

tierunassign
Initiates the unassigning volume action on the specified volumes.

etmonpause
Specifies that Easy Tier monitoring of this volume is paused. During the pause, all Easy Tier storage migrations are unaffected, but no new migration plans are formed.

etmonresume
Specifies that Easy Tier monitoring of this volume is resumed. All Easy Tier storage migrations are unaffected.

etmonreset
Specifies that all Easy Tier monitoring data (history), including migration plans are erased. All new plans are based on new monitoring data.

mksafeguardedcap
Creates the safeguarded backup capacity for one or more of the specified volumes.

rmsafeguardedcap
Deletes the safeguarded backup capacity, including all safeguarded copies, that were defined for one or more of the specified volumes.

migsafeguardedcap
Migrates the safeguarded backup capacity to a new pool for one or more of the specified volumes.

migsafeguardedcapcancel
Cancels migrating the safeguarded capacity to a new pool for one or more of the specified volumes.

migsafeguardedcappause
Pauses migrating the safeguarded backup capacity to a new pool for one or more of the specified volumes.

migsafeguardedcapresume
Resumes migrating the safeguarded backup capacity to a new pool for one or more of the specified volumes.

-eam
(Optional) Specifies the extent allocation method as follows:

rotateexts
Specifies that the extents for each new logical volume are allocated across all available ranks, and is also known as storage-pool striping. This value is the default.

rotatevols
Specifies that the extents for each new logical volume are allocated from each successive rank. This means that the extents for a particular volume are allocated from one rank, while the extents for the next volume are allocated from the next successive rank, and so on.

Note: You can specify only the **-eam** parameter if **-action** migstart is also specified.

-extpool *extent_pool_ID*
(Optional) Specifies the multiplier value. Selecting the correct backup capacity multiplier depends on a number of factors. For example, the provisioned capacity being backed up, the number of copies, the interval over which the copies are kept, and the predicted destagerate of the volume data.

Note: When the command returns, the volume migration might still be occurring. It can be available for I/O and copy services during migration. Its configstate can indicate that it is migrating.

-multiplier *multiplier*
(Optional) Specifies the multiplier value. Selecting the correct backup capacity multiplier depends on the provisioned capacity being backed up, the number of copies, the interval over which the copies are kept, and the predicted destagerate of the volume data.

-cap *capacity*
(Optional) Specifies the safeguarded backup capacity for one or more specified volumes. The unit type is set with the parameter **-captype**, if the unit type is not set. The default is *cylinders*.

-captype cyl | mod1
(Optional) Specifies the unit type for safeguarded backup capacity. If the unit type is not set, the default is *cylinders*.

cyl Specifies the capacity unit type as cylinders. This capacity unit type can be selected only for CKD storage. One cylinder = 15 tracks.

mod1 Specifies the capacity unit type in mod1 units. This capacity unit type can be selected only for CKD storage (mod1 = 1113 cylinders).

-tier flashtier0 | flashtier1 | flashtier2 | ent | nl | nlexclude | hddexclude
(Optional) Specifies which tier the volume is assigned to. This option is required with the **-action tierassign** parameter.

flashtier0
The tier of high-performance flash drives.

flashtier1
The first tier of high capacity flash drives.

flashtier2
The second tier of high capacity flash drives.

ent Enterprise tier, which consists of drives with speeds of 10K RPM, 15K RPM, or a mixture of 10K RPM and 15K RPM speeds.

nl Nearline tier consists of high-volume disks that are either SATA or SAS Nearline drives.

nlexclude
Enterprise tiers but not a Nearline tier.

hddexclude
Any tier except enterprise (ENT) or nearline (NL).

" - " No assign action was specified.

-assignorder etdata|access
(Optional) Specifies the order in which the data is migrated. This option is valid only with the **-action tierassign** parameter.

etdata While all data is scheduled to migrate, the migration order is based on the prioritization of the data as specified in the Easy Tier heat map. This value allows the specified volume to be pre-staged onto the specified tier. The value "etdata" is the default value if **-assignorder** is not specified.

access While all data is scheduled to migrate, the data is migrated only when accessed. In other words, data that is never accessed is never migrated to the specified tier.

-duration *time*
(Optional) Specifies the hours of the pause time in ISO 8601 format. For example, **-duration** 24H. The maximum value of the time is a week, which is 168 hours (168H). You can specify this option only with **-action etmigpause** or **etmonpause** parameters.

Note: If you want the duration of the pause to be infinite, you must specify **-duration** 0H. Otherwise, if you do not specify a value with the **-duration** parameter, the default is 168H.

volume_ID ... | **-**
(Required) Specifies an array of one or more CKD base volume IDs or volume ID ranges to modify.

A volume ID range is defined by two volume IDs that are separated by a dash. Multiple volume IDs or volume ID ranges must be separated with a blank space between each ID.

Example: 0100-010F 0180-018F 0120

The volume ID format is four hexadecimal characters *LLVV that represents* the following values:

LL Specifies the logical control unit number, 00 - FE

VV Specifies the volume number, 00 - FF

You must fully qualify the volume ID with manufacturer, machine type, and serial number if you do not use the **-dev** parameter.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) if you are in the DS CLI interactive mode.

Example

Invoking the manageckdvol command

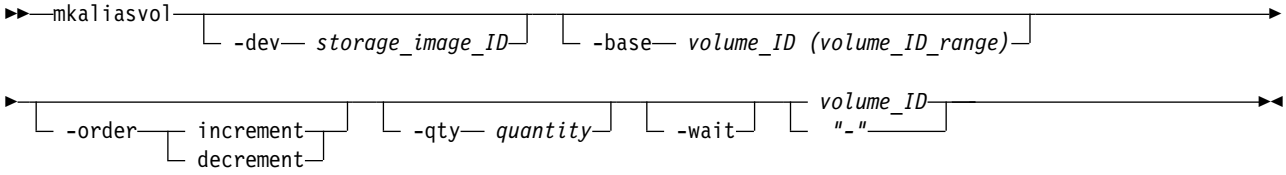
```
dscli> manageckdvol -dev IBM.2107-75FA120
-action migstart -extpool P2 0100
```

The resulting output

```
CMUC00000I manageckdvol:
CKD Volume 0100 action migstart executed successfully.
```

mkaliasvol

The **mkaliasvol** command creates IBM Z CKD alias volumes (referred to as parallel access volumes or PAVs) in a storage image.



Parameters

Note: Volumes are automatically assigned to the FICON/ESCON - ALL volume group ID 10.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-base *volume_ID (volume_ID_range)*
(Required) Specifies an existing base CKD volume ID or a volume ID range.

Note: You cannot use multiple volume IDs separated by commas and multiple ID ranges in combination. This combination is rejected.

Use the **-base** parameter to create one or more CKD alias volumes that are assigned to the specified base CKD volume ID. The LCU ID component for all volume IDs must be identical.

The alias volume IDs are assigned consecutively in the order specified by the **-order** parameter. The following examples show the processing affects of the **-order** parameter:

```
dscli> mkaliasvol -base 0000 -order increment -qty 2 0080
creates two alias volumes 0080 and 0081 for base volume 0000.
```

```
dscli> mkaliasvol -base 0000-003F -order increment -qty 2 0080
creates two alias volumes for each base volume as follows:
0080,0081 for base volume 0000
0082,0083 for base volume 0001
...
00FE,00FF for base volume 003F
```

-order increment | decrement
(Optional) Specifies the order in which alias volume IDs are assigned. For example:

```
dscli> mkaliasvol -base 0000-003F -order decrement -qty 2 00FF
creates two alias volumes for each base volume as follows:
00FF,00FE for base volume 0000
00FD,00FC for base volume 0001
...
0081,0080 for base volume 003F
```

Note: If the **-order** parameter is not specified the default value is decrement.

-qty *quantity*
(Optional) Specifies the number of alias volumes that are assigned to each specified CKD base volume.

If you do not specify the **-qty** parameter, then one alias volume is created for each base volume specified.

If you specify the **-qty** parameter, you have to indicate the number of alias volumes that you want to assign to each specified CKD base volume.

Note: The total number of base volumes plus the number of alias volumes must be less than or equal to 256.

-wait
(Optional) Delays the command response until the volume configuration processes complete.

volume_ID -
(Required) Identifies the starting alias volume ID in a sequence of volume IDs to be created

The volume ID format is four hexadecimal characters *LLVV* that represent the following values:

LL Specifies the logical control unit number, 00 - FE

VV Specifies the volume number, 00 - FF

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the mkaliasvol command

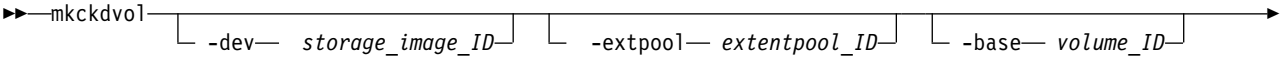
```
dscli> mkaliasvol -dev IBM.2107-75FA120 -base 0100-010F
-order decrement -qty 2 01FF
```

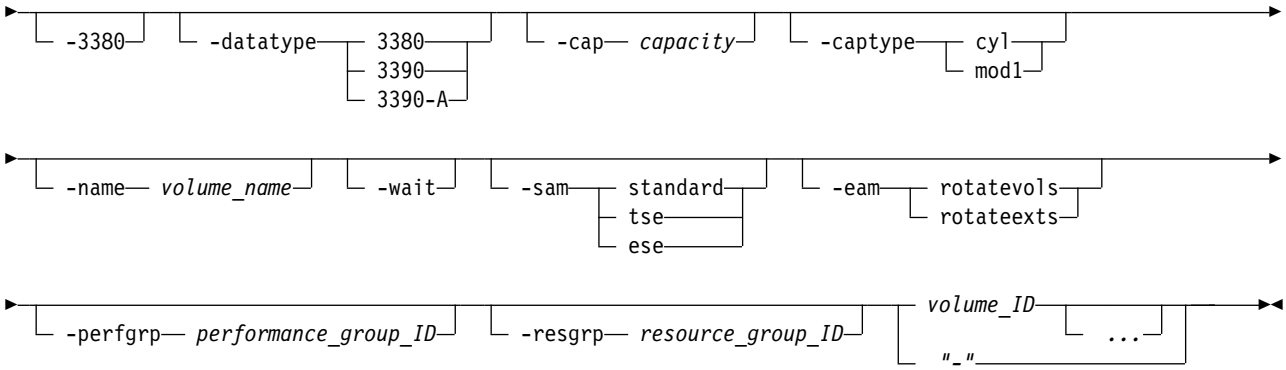
The resulting output

```
CKD Volume 01FF successfully created.
CKD Volume 01FE successfully created.
...
CKD Volume 01E1 successfully created.
CKD Volume 01E0 successfully created.
```

mkckdvol

The **mkckdvol** command creates IBM Z count key data (CKD) base or CKD alias volumes in a storage image.





Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The **-dev** parameter temporarily overrides any defined value for *devid* for the current command.The storage image ID is required for the following situations:

- If you do not specify a fully qualified extent pool ID
- If you do not set the *devid* variable in your profile or through the **setenv** command
- If the HMC is aware of more than one storage image

-extpool *extentpool_ID*
(Optional) Creates the base or alias volumes from data extents that are contained in this extent pool. The extent pool storage type defines the volume storage type. An extent pool ID is a four-digit decimal number with no leading zeros, prefixed with the letter *P*.

Note: This parameter is ignored if the **-base** parameter is specified.

-base *volume_ID*
(Optional) Specifies an existing base CKD volume ID. The volume ID format is four hexadecimal characters *LLVV*, that represents the following values:

LL Specifies the logical control unit number, 00 - FE

VV Specifies the volume number, 00 - FF

Use the **-base** parameter to create one or more CKD alias volumes that are assigned to the specified base CKD volume ID. The LCU ID component for all volume IDs must be identical.

Note: Ensure that you use the **mkaliasvol** command rather than **mkckdvol** command to create alias volumes.

-3380
(Optional) This parameter is equivalent to the **-datatype** *3380* parameter and it is recommended that you use the **-datatype** parameter if you need to specify the data type.

-datatype 3380|3390|3390-A
(Optional) Specifies the volume data type.

For Extended Addressing Volumes (EAV), which have a capacity greater than 65520 cylinders, the data type 3390-A must be specified. You can also specify the data type 3390-A for volumes smaller than 65520 cylinders.

If the **-datatype** or the **-3380** parameters are *not* specified, then the default data type is 3390-A for capacities greater than 65520 cylinders, and 3390 for smaller capacities.

Notes:

1. You can specify the data type 3390–A for DS8000 models only.
2. You must ensure that the volume data type is compatible with the host systems that access this volume.
3. The **-datatype** parameter is ignored when the **-base** parameter is specified.
4. You cannot use the **-datatype** parameter and the **-3380** parameter in the same command.

-cap *capacity*
(Optional) Specifies the quantity of CKD cylinders that are allocated to this volume. The capacity can also be specified in mod1 units by using the **-captype** parameter.

- For 3380 volumes, the *capacity* value can be 2226 (1.59 GiB) or 3339 (2.37 GiB).
- For 3390 volumes, capacity can be specified in increments of 1 in the range 1 - 65520 (849KiB to 55.68 GiB).
- For 3390–A volumes, for values less than 65520 cylinders, capacity is specified in increments of 1. Capacities greater than 65520 cylinders are specified in increments of 1113. The maximum volume size varies and depends on the DS8000 model and type.

Note: This parameter is ignored if the **-base** parameter is specified, and this parameter is required if the **-base** parameter is not specified.

-captype cyl | mod1
(Optional) Specifies the unit type of the capacity that is given by using the **-cap** parameter. The default is cyl. A mod1 unit is equivalent to 1113 cylinders, and 1263.28 cylinders is equivalent to 1 GiB.

-name *volume_name*
(Optional) Specifies your nickname for the CKD base volumes that are created by this command. Your volume name cannot exceed 16 characters. It can contain one of the following wildcards:

- (#d) insert volume ID (decimal)
- (#h) insert volume ID (hexadecimal)

Note: The **-name** parameter is ignored when the **-base** parameter is specified.

-wait
(Optional) Specifies that the command response is delayed until the volume configuration processes complete.

-sam standard | tse | ese
(Optional) Specifies the storage allocation method as follows:

standard
Designates that the system fully allocates the volume with real extents at volume creation time. This value is the default.

tse Designates that a track space-efficient logical volume contains a set of virtual extents that are associated with the space-efficient storage in the same extent pool. The physical space for a given logical track on a track space-efficient logical volume is dynamically allocated and deallocated from the repository in the space-efficient storage.

Note: Track space-efficient (TSE) repositories are required to create TSE volumes. However, starting with DS8000 Release 8, TSE repositories can no longer be created.

ese Designates that an extent space efficient (ESE) logical volume is provisioned with a set of virtual extents that are associated with the space efficient storage in the same extent pool. Physical space for an extent space efficient logical volume is dynamically allocated and de-allocated from the extent pool. ESE volumes are used for thin provisioning.

-eam rotatevols | rotateexts
(Optional) Specifies the extent allocation method as follows:

rotateexts
Specifies that the extents for each new logical volume are allocated across all available ranks, and is also known as storage-pool striping. This value is the default.

rotatevols
Specifies that the extents for each new logical volume are allocated from each successive rank. This means that the extents for a particular volume will be allocated from one rank, while the extents for the next volume will be allocated from the next successive rank, and so on.

-perfgrp *performance_group_ID*
(Optional) Specifies the performance group ID that the volumes are assigned to. The performance group ID begins with the letters *PG*. The default is *PG0*.

-resgrp *resource_group_ID*
(Optional) Specifies the resource group that the volumes are assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number. The default is *RG0*.

volume_ID ... | -
(Required) Specifies an array of one or more CKD base or alias volume IDs or volume ID ranges to be created. The volume IDs must share a common logical control unit ID.

Note: Volumes are automatically assigned to the FICON/ESCON – ALL volume group ID 10.

The volume ID format is four hexadecimal characters *LLVV*, that represents the following values:

LL Specifies the logical control unit number, 00 - FE

VV Specifies the volume number, 00 - FF

A volume ID range is defined by two volume IDs that are separated by a dash.

You must separate multiple volume IDs or ranges of volume IDs with a blank space between each ID or range of IDs.

Note: Multiple volumes can be created with a single request, but all volumes must have the same capacity, extent pool, and data type.

Example: 0100-010F 0180-018F 0120

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) if you are in the DS CLI interactive mode.

Example

Invoking the mkckdvol command

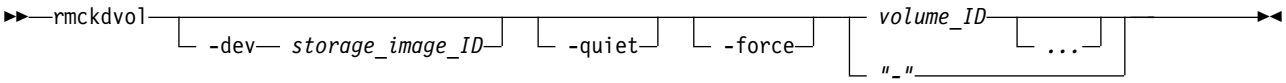
```
dscli> mkckdvol -dev IBM.2107-75FA120
-extpool P1 -name my_volume_#d -cap 3339
-perfgrp PG17 -sam ese 0100 0101 0102 0103
```

The resulting output

```
CKD volume 0100 successfully created.
CKD volume 0101 successfully created.
CKD volume 0102 successfully created.
CKD volume 0103 successfully created.
```

rmckdvol

The **rmckdvol** command deletes count key data (CKD) base or alias volumes from a storage image.



Parameters

Note: Before Release 5.1, the DS8000 system did not check before deleting a volume. With Release 5.1 and later, the DS8000 system does not delete a volume that is in use. The phrase *in use* means that the volume is participating in a Copy Services relationship or is in a z/OS path group. Use the **-force** parameter to bypass the in-use checking and delete the volume.

A specific use of this command is made when you are confronted with a volume or volumes that are in a configuration state of "configuration error." To correct this configuration state, issue the **rmckdvol** command for each affected volume. You can specify a volume range according to the command specifications when it is appropriate.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for all logical volumes, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-quiet
(Optional) Turns off the volume removal confirmation prompt for this command.

-force
(Optional) Specifies that you want to remove volumes without checking to see whether the volumes are in use. The phrase *in use* means that the volume is participating in a Copy Services relationship or is in a z/OS path group. If multiple volumes are specified and some are in use and some are not, the ones not in use are deleted.

volume_ID ... -
(Required) An array of one or more CKD base or CKD alias volume IDs or volume ID ranges to be removed. Accepts a fully qualified volume ID, which includes the storage image ID or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened volume ID format is four hexadecimal characters *LLVV*, that represents the following values:

LL Specifies the logical control unit number, 00 - FE

VV Specifies the volume number, 00 - FF that is contained by a logical control unit (logical subsystem).

Note: An alias volume that is associated with a CKD base volume is automatically deleted before deletion of the CKD base volume.

A volume ID range is defined by two volume IDs that are separated by a dash.

You must separate multiple volume IDs or ranges of volume IDs with a blank space between each ID or range of IDs.

Example: 0100-010F 0180-FEFF 0120

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmckdvol command

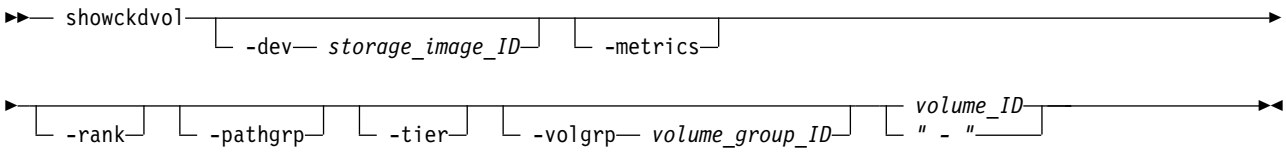
```
dscli> rmckdvol -dev IBM.2107-75FA120 0000 0001
```

The resulting output

```
The alias volumes associated with a CKD base volume are
automatically deleted before deletion of the CKD base volume.
Are you sure you want to delete CKD volume 0000? [y/n]
Are you sure you want to delete CKD volume 0001? [y/n]
```

showckdvol

The **showckdvol** command displays detailed properties of an individual count key data (CKD) volume. Detailed information includes performance metrics of a CKD volume and information about Safeguarded Copy capacity.



Parameters

Note: All performance counts are an accumulation from the most recent counter wrap or counter reset. A reset of the volume performance counters occurs in association with the following events:

- The storage system is turned on.
- A server failure occurred, and the server failover or failback sequence was initiated.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of the manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID. It is also required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-metrics
(Optional) Displays the volume ID and performance metrics for the specified volume.

Notes:

1. All performance counts are an accumulation since the most recent counter wrap or counter reset. Volume performance counters are reset on a power-up sequence. Volume performance counters are reset by a server failover and failback sequence.
2. Do not use the **-metrics** parameter with the **-pathgrp**, **-rank**, **-tier**, or **-volgrp** parameters.

-pathgrp
(Optional) Displays the path group status table, which contains path group information. This information includes the path group ID, the grouped, reserved, and path mode status.

Do not use the **-pathgrp** parameter with the **-metrics**, **–rank**, **-tier**, or **-volgrp** parameters.

-rank
(Optional) Specifies that a rank extents table is to be displayed. This table displays the set of ranks that the logical volume has extents that are configured on and the number of extents for that logical volume.

Do not use the **-rank** parameter with the **-metrics**, **-pathgrp**, **-tier**, or **-volgrp** parameters.

-tier
(Optional) Displays the tier distribution table. The table lists the set of tiers that have storage that is allocated for the specified logical volume and the percentage of the logical volume that is allocated on each tier.

Do not use the **-tier** parameter with the **-metrics**, **-pathgrp**, **-rank**, or **-volgrp** parameters.

-volgrp *volume_group_ID*
(Required if you do not specify the *volume_ID* parameter.) Specifies that the CKD volumes that are associated with the designated volume group ID is to be displayed. Only one volume group for CKD volumes exists and it contains all volumes.

Notes:

- The **-volgrp** parameter can be used only when you are doing a query for performance metrics.
- Do not use the **-volgrp** parameter with the *volume_ID* parameter.
- Do not use the **-volgrp** parameter with the **-metrics**, **-pathgrp**, **-rank**, or **-tier** parameters.

volume_ID **-**
(Optional) Specifies that you want the system to display detail information about the designated volume. The volume ID format is four hexadecimal characters *LLVV* that represents the following values:

LL Specifies the logical control unit number, 00 - FE.

VV Specifies the volume number, 00 - FF.

If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output reports that are associated with the **showckdvol** command by using the **-rank** parameter. When the rank parameter is specified, a rank extents table is also displayed. It appears at the end of the regular report.

Invoking the showckdvol to show volume properties

Note: The following example is based on the use of the **showckdvol** command for a 3390 volume with 3339 cylinders. When the rank parameter is specified, a rank extents table is displayed at the end of the regular report.

```
dscli> showckdvol -dev IBM.2107-1300861 -rank 1410
```

The resulting output

Name	ID	acc state	data state	config state	device MTM	volser	data type
My_ volume_ 1410	1410	Online	Normal	Normal	3390-3	A03967	3390

voltype	orgb-vols	addrgrp	extpool	exts	cap (cyl)	cap (2^30B)	cap (10^9B)	Ranks
CKD Base	–	1	P2	3	3339	2.6	2.8	2

sam	repcapalloc	eam	reqcap (cyl)	cap (Mod1)
ESE	1.7	-	3339	0.0

realextents	virtualextents	realcap (cyl)	migrating cap(cyl)	migrating	migratingfrom	perfgrp	resgrp
1	0	0	-	PG0	RG0		

tierassignstatus	tierassignerror	tierassignorder	tierassigntarget	%tierassigned
Assigning	-	ETdata	FlashTier0	54

safeguardedcap (cyl)	safeguardedloc	usedsafeguardedcap cyl)	safeguardeded
1234	mypool_1	0	yes

```
============================Rank extents============================
```

Rank	Extents	Rank Extents Capacity (cyl)
R0	1	1113
R2	2	1113

Report field definitions (*-metrics* parameter not specified)

Name
Indicates the nickname that you assigned for this volume object.

ID Indicates the unique identifier that is assigned to this volume object.

Accstate
One of the following values can be displayed:

Online
Indicates that the logical volume is accessible to a host.

Fenced
Indicates that the logical volume is in the volume fenced state and is not accessible to the host.

" - "
Indicates that the logical volume is designated as a CKD alias volume.

Datastate
One of the following values can be displayed:

Normal
Indicates that none of the other data states apply. The access state is Online.

Pinned
Indicates that none of the other data states apply and the logical volume has one or more pinned non-retryable tracks. The access state is Online.

Read only
Indicates that the logical volume is read-only because one or more extents on the logical volume are on a rank in the read-only data state. The access state is Online.

Inaccessible
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the inaccessible data state. The access state is Fenced.

Virtual space fault
Indicates that the logical volume has a space-efficient storage allocation method, and enough space was not available to convert a virtual logical track to a real logical track. The access state is Online.

Indeterminate data loss
Indicates that the following data states do not apply and that one of the following conditions occurred:

Data states that do not apply:

- Rank failed
- Rank repairing
- Rank repaired
- Global inaccessible
- Global lost data

Conditions - One of the following conditions occurred:

- Committed write data was lost before it was destaged and the track identifiers that are associated with the data are unknown.
- Data was lost that indicated extents on the logical volume were active FlashCopy targets.

The access state is Fenced.

Rank failed
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the failed data state. The access state is Fenced.

Rank repairing
Indicates that one or more extents that are associated with the logical volume are on ranks that are in the repairing data state. The access state is Fenced.

Rank repaired
Indicates that one or more extents that are associated with the logical volume are on ranks that were in the repairing state, but are not in the repairing state now. The access state is Fenced.

Global inaccessible
Indicates that the global metadata that is associated with the logical volume configuration is inaccessible. Some of the data that is associated with the logical volume might be inaccurate. The access state is Fenced.

Global lost data
Indicates that global metadata that is associated with the logical volume configuration was lost. As a result, some of the data that is associated with the logical volume might be inaccurate. The access state is Fenced.

NVS data inaccessible
Indicates that active NVS data is inaccessible for one or more logical volumes of an LSS group. The logical volumes in the LSS group cannot be made accessible. The access state is fenced.

Indicates that the logical volume is designated as a CKD alias.

Configstate
One of the following values can be displayed:

Normal
Indicates that no logical volume configuration operations are in progress, and the volume is not being deconfigured, merged, or migrated.

Configuring
Indicates that the logical volume is being configured for the first time.

Reconfiguring
Indicates that the logical volume is allocating or deallocating extents due to a modification of the requested capacity attribute after initial creation.

Deconfiguring
Indicates that the logical volume is starting to be deleted.

Configuration error
Indicates that the initial configuration did not complete successfully. This state reflects an internal error condition and not an error in the request to create the volume. If you have a volume in this state, use the **rmfbvol** command to delete each volume that is listed with the configuration state of "configuration error".

Merging
Indicates that the volume is in the process of merging. For example, merging from one extent pool to a different extent pool.

Migrating
Indicates that the volume is starting to be migrated, or waiting to be migrated.

Migration Cancelled
Indicates that the volume was starting to be migrated and then the 'migcancel' action of the **manageckdvol** command was entered. This process left some of the extents waiting to be migrated in the source pool while other extents were already migrated to the target pool. Migration stopped, and cannot be resumed. If you have a volume in this state, try to migrate it again to the original source or target extent pool.

Migration Paused
Indicates that the volume was starting to be migrated and then the 'migpause' action of the **manageckdvol** command was entered. Migration stopped, but can be resumed.

Migration Error
Indicates that the volume migration process failed to complete successfully. This state reflects an internal error condition and not an error in the user's request to migrate a volume. If you have a volume in this state, try to migrate it again to the original source or target extent pool.

Reconfiguration error
Indicates that the reconfiguration request did not complete successfully.

Deconfiguration error
Indicates that a request to delete a volume did not complete successfully. This state reflects an internal error condition and not an error in the request to remove the volume. To correct this state, you must reissue the **rmfbvol** command for the designated volume.

Transposition Error
Indicates that the volume is in an extent pool that was unsuccessfully merged. This state reflects an internal error condition. *Corrective action:* Use the **chextpool** command with the **-merge** parameter again to redrive the merge extent pool and to correct this state.

deviceMTM
One of the following device MTMs can be displayed:

- 3380-2
- 3380-3
- 3390-3
- 3390-9

Volser
Indicates the volume serial number. A volume serial number is 6 bytes of data, displayed as 6 characters.

Datatype
Indicates the volume data type setting (3380 or 3390).

Voltype
Indicates that the logical volume is one of the following types:

- CKD base
- CKD alias

Orgbvols (original base vol)
One of the following original base volumes can be specified:

- Indicates the original base CKD volume ID to which this CKD alias volume is assigned.

  Note: For a CKD Alias type volume, the base and alias volume IDs share a common LCU ID.
- " - " is displayed for a CKD base volume type.

Addrgrp
Indicates the address group that contains this volume object. An address group ID is one hexadecimal character (0 - F).

Extpool
Indicates the extent pool ID.

Note: Volumes that belong to an encrypted extent pool are encrypted. You can see the key group of an extent pool by using the **lsextpool -l**, or **showextpool** commands.

Exts
Indicates the number of real and virtual extents that are used by the designated volume ID.

Cap(cyl)
Indicates the quantity of volume cylinders that are available for host system access.

Cap (2^30B)
Indicates the size of volume that is available for host system access in gibibytes (GiB).

Cap (10^9B)
Indicates the size of volume that is available for host system access in decimal gigabyte (GB) units.

Ranks
Indicates the number of ranks the volume on.

SAM
Indicates the storage allocation method. The following values are displayed:

standard
Designates that the system fully allocated the volume with real extents at volume creation time.

tse
Designates that a track space-efficient logical volume contains a set of virtual extents that are associated with the space-efficient storage in the same extent pool. Physical space for a given logical track on a track space-efficient logical volume is dynamically allocated and deallocated from the repository in the space-efficient storage.

Note: Track space-efficient (TSE) repositories are required to create TSE volumes. However, starting with DS8000 Release 8, TSE repositories can no longer be created.

ese
Designates that an extent space efficient logical volume is provisioned with a set of virtual extents that are associated with the space efficient storage in the same extent pool. Physical space for an extent space efficient logical volume is dynamically allocated and de-allocated from the extent pool.

Repcapalloc
Indicates the allocated physical repository capacity of the track space-efficient storage. This value is calculated on the available repository capacity as a result of writes to the track space-efficient volume. This value is displayed in the format of X.Y, where X is in whole gibibytes (1 GiB = 2^30 B). Y represents tenths of a GiB, which is limited to a single digit (0 - 9).

Note: A " - " value is displayed in this column if the value displayed in the SAM column is Standard or ESE.

EAM
Indicates the extent allocation method that is used if the volume is migrated or expanded. One of the following values is displayed:

legacy
Indicates that the volume was created before the use of the current algorithm.

rotateexts
Indicates that the extents for each new logical volume are allocated across all available ranks, and is also known as storage-pool striping. This value is the default.

rotatevols
Indicates that the extents for each new logical volume are allocated from each successive rank. This approach means that the extents for a particular volume is allocated from one rank. The extents for the next volume will be allocated from the next successive rank, and so on.

managed
Indicates that the extents are currently managed by Easy Tier, and the extents for any new volumes are initially allocated across all available ranks in the lowest tier of storage.

" – "
A " - " value is displayed if the extent allocation method does not apply, for example, track space-efficient logical volumes.

Reqcap (cyl)
Indicates the requested quantity of volume CKD cylinders (for example, 3339).

cap (Mod1)
Indicates the quantity of Mod1 units (Mod1 = 1113 cylinders).

realextents
Indicates the number of real extents that are used by the logical volume.

virtualextents
Indicates the number of virtual extents that are used by the logical volume.

realcap (cyl)
Indicates the number of cylinders used by the specified logical volume.

migratingcap (cyl)
The number of cylinders for the specified volume that are currently being migrated.

migrating
The number of extents for the specified volume that are currently being migrated.

migratingfrom
A list of one or more extent pool IDs where the extents are migrating from. If no migrating extents exist, a dash "-" is displayed. `Unknown` Is displayed if information about the extent pool IDs is not available.

perfgrp
Indicates the performance group ID that the volume is assigned to. The performance group ID begins with the letters *PG* and ends with a decimal number.

resgrp
Indicates the resource group ID that the volume is assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

tierassignstatus
Status of assigning a volume to a target tier.

Assign Pending
An assign action was specified, but has not started.

Assign Pending Hardware
An assign action was specified, but paused because of a hardware condition.

Assigning
An assign action is in progress.

Assigned
The volume was assigned to the specified tier.

Unassign Pending
An unassign action was specified, but has not started.

Error An assign action that failed. See the **tierassignerror** value for the reason.

Unknown
An assign action was specified, but the specific action is unknown.

" – " No assign action was specified (none).

tierassignerror
Failure reason if the status is Error.

Easy Tier not active
Easy Tier is not active. See the `etmanaged` column from **lsextpool** to see whether the volume is in a pool that is managed by Easy Tier.

Use `manageckdvol -action tierunassign` to unassign the volume, ensure that the pool is managed by Easy Tier (see **chsi**), and then use `manageckdvol -action tierassign` to assign the volume again.

Target Tier not available
The specified target tier does not currently exist. Use `manageckdvol -action tierunassign` to unassign the volume, ensure that space is available on the specified tier, and then use `manageckdvol -action tierassign` to assign the volume again.

Tier definitions have changed
The target tier was valid, but defined tiers changed internally and the target tier is no longer valid. Use `manageckdvol -action tierunassign` to unassign the volume, and then use `manageckdvol -action tierassign` to assign the volume again.

" – " The assign action status is not "Error".

tierassignorder
Method that is used to define the assigning order.

Access
Assign extents only when accessed.

ETdata
Assign high usage extents first, based on East Tier data.

Unknown
Unknown assigning order method.

" – " No assigning action was specified.

tierassigntarget
Assigning action target tier.

FlashTier0
The tier of high performance flash drives.

FlashTier1
The first tier of high capacity flash drives.

FlashTier2
The second tier of high capacity flash drives.

ENT Enterprise tier consists of drives with speeds of 10K RPM, 15K RPM, or a combination of drives of both speeds.

NL Nearline (NL) tier consists of high volume SATA or SAS Nearline disk drives.

HDDExclude
Any tier except enterprise (ENT) or nearline (NL).

NLExcluded
Enterprise tiers but not a Nearline tier.

Unknown
Assigning action was specified, but the target tier is unknown.

" – " No assigning action was specified.

%tierassigned
Percentage of the volume that is assigned. The value is 0 (zero) if the volume is not assigned to a tier.

etmonpauseremain
Specifies the pause in Easy Tier monitoring. One of the following values is displayed:

0H1M-168H0M
Specifies the time (in hours and minutes) that remains of the pause in the Easy Tier monitoring process.

infinite
Specifies that Easy Tier monitoring remains paused until a resume action is submitted.

- A dash (-) specifies that Easy Tier monitoring is not paused.

unknown
Specifies that the system failed to query the time that remains of the pause.

etmonitorreset
Easy Tier extent pool monitoring reset date is as follows:

date
Specifies the date of the last Easy Tier monitoring reset in ISO 8601 format: `yyyy-MM-dd'T'HH:mm:ssZ`, where

yyyy The year

MM The month (01-12)

dd The day (01-31)

T The single letter T without quotes

HH The hour (00-23)

mm The minutes (00-59)

ss The seconds (00-59)

Z The time zone offset from UTC [-HHmm | +HHmm]

unknown
Specifies that the date in which Easy Tier monitoring of this extent pool was last reset is not known.

unsupported
Specifies that Easy Tier extent pool management is not supported.

safeguardedcap(cyl)
Specifies the amount of safeguarded backup capacity that is reserved for safeguarded copies.

safeguardedloc
Specifies the pool where the safeguarded copies are written to.

usedsafeguardedcap(cyl)
Specifies the amount of physical pool capacity that is used by the safeguarded copies.

safeguardeded
Specifies whether Safeguarded Copy is enabled for the specified volume.

Report field definitions (*-rank* parameter specified)

Rank (Rank Extent table)
Indicates the rank ID.

Extents (Rank Extents table)
Indicates the number of extents for the volume on the rank.

Rank Extents Capacity (cyl) (Rank Extents table)
Indicates the number of cylinders for the volume on the rank.

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output reports that are associated with the **showckdvol** command by using the **-metrics** parameter.

Invoking the showckdvol command to show performance metrics

```
dscli> showckdvol -metrics IBM.2107-75FA120/0101
```

The resulting output

ID	Date	norm rdrqts	norm rdhits	norm write req	norm write hits	seq read reqs	seq read hits	seq write req
IBM.2107-75FA120/0101	10/11 /04 02:23:49	10000	10000	10000	10000	10000	10000	10000

seqwrite-hits	cachfwr-reqs	cachfwr-hits	cachfw-reqs	cachfw-hits	inbcach-load	bypass-cach	seq DASD trans
10000	10000	10000	10000	10000	10000	10000	10000

DASD-trans	cache-trans	NVS-spadel	norm write ops	seqwrite-ops	rec cache mis	qwrite-prots	CKDir-trkac
10000	10000	10000	10000	10000	10000	10000	10000

CKD irtrk hits	cachsp-delay	timelow-ifact	phread	phwrite	phwrite	phbyte-read	phbyte-writ
10000	10000	10000	10000	10000	10000	10000	10000

recmo-reads	sfile trk reads	contam-wrts	PPRC-trks	NVS-spallo	timeph-read	timeph-write	byte-read
10000	10000	10000	10000	10000	10000	10000	10000

bytewrit	timeread	timewrite	zHPFRead	zHPFWrite
10000	10000	10000	-	-

zHPFPrefetchReq	zHPFPrefetchHit
0	0

GMCollisionsSidefileCount	GMCollisionsSendSyncCount
0	0

Report field definitions (*-metrics* parameter specified)

ID Indicates the unique identifier that is assigned to this volume object.

Date
Indicates the current time stamp for the volume performance counters.

normrdrqts
Indicates the number of normal read operations that are assigned by a host to a volume.

normrdhits
Indicates the number of normal read operations where data did not move to or from a storage device.

normwritereq
Indicates Write Normal I/O Requests

normwritehits
Indicates DASD Fast Write I/O Request Hits

seqreadreqs
Indicates Search/Read Sequential I/O Requests

seqreadhits
Indicates Search/Read Sequential I/O Request Hits

seqwritereq
Indicates Write Sequential I/O Requests

seqwritehits
Indicates DASD Fast Write Sequential I/O Request Hits

cachfwrreqs
Indicates Search/Read Cache Fast Write I/O Requests

cachfwrhits
Indicates Search/Read Cache Fast Write I/O Request Hits

cachfwreqs
Indicates Cache Fast Write I/O Requests

cachfwhits
Indicates Cache Fast Write I/O Requests Hits

inbcachload
Indicates Inhibit Cache Loading I/O Requests that operate with DASD

bypasscach
Indicates Bypass Cache I/O Requests

seqDASDtrans
Indicates Sequential DASD to Cache Transfer Operations

DASDtrans
Indicates DASD to Cache Transfer Operation Count

cachetrans
Indicates Cache to DASD Transfer Operation Count

NVSspadel
Indicates DASD Fast Write Operations Delayed Due to nonvolatile storage Space Constraints

normwriteops
Indicates Normal 'DASD Fast Write' Write Operation Counts

seqwriteops
Indicates Sequential Access 'DASD Fast Write' Write Operation Counts

reccachemis
Indicates Number of record cache Read Misses

qwriteprots
Indicates Quick Write Promotes

CKDirtrkac
Indicates Irregular Track Accesses

CKDirtrkhits
Indicates Irregular Track Accesses Hits

cachspdelay
Indicates Operations Delayed Due To Cache Space Constraints

timelowifact
Indicates Milliseconds of lower interface I/O activity for the indicated device.

phread
Indicates Physical Storage Read Operations

phwrite
Indicates Physical Storage Write Operations

phbyteread
Indicates Physical Storage Bytes Read in 128 KB increments.

phbytewrit
Indicates Physical Storage Bytes Written in 128 KB increments.

recmoreads
Indicates Record Mode Read Operations

sfiletrkreads
Indicates the Number of tracks that are read from the Concurrent Copy or XRC Sidefile.

contamwrts
Indicates the Number of Contaminating writes for a Concurrent Copy or XRC volume

PPRCtrks
Indicates the Number of tracks or portion of tracks that were transferred to the secondary device of a PPRC pair.

NVSspallo
Indicates the NVS Space Allocations

timephread
Indicates the physical storage read response time in 16 ms increments.

timephwrite
Indicates the physical storage write response time in 16 ms increments.

byteread
Indicates the number of bytes that are read in 128 KB increments

bytewrit
Indicates the number of bytes that are written in 128 KB increments.

timeread
Indicates the accumulated response time for all read operations.

timewrite
Indicates the accumulated response time for all write operations.

zHPFRead
Indicates the High Performance FICON (HPF) Read I/O Requests for volume performance statistics.

zHPFWrite
Indicates the HPF Write I/O Requests for volume performance statistics.

zHPFPrefetchReq
Indicates the number of HPF Pre-fetch I/O requests.

zHPFPrefetchHit
Indicates the number of HPF Pre-fetch I/O request hits.

GMCollisionsSidefileCount
Indicates the number of Global Mirror Collisions sidefile.

GMCollisionsSendSyncCount
Indicates the number of Global Mirror Collisions Send Synchronous Count.

Example

Specifying the -tier parameter

If you specify the **-tier** parameter, a tier distribution table is appended to the end of the display.

```
dscli> showckdvol –tier 0000
```

The resulting output

```
Name             myvol0800
ID               0800
accstate         Online
datastate        Normal
```

```
configstate       Normal
...
migrating         20
perfgrp           PG0
migratingfrom     P0
resgrp            RG1
tierassignstatus  Assigning
tierassignerror   -
tierassignorder   ETdata
tierassigntarget  FlashTier0
%tierassigned     54
etmonpauseremain  1H44M
etmonitorreset    2013-07-26T14:00:00+07
============== Tier Distribution =================
Tier %allocated
==============
FlashTier0     54
ENT            46
```

Report field definitions (*-tier* parameter is specified)

Tier Tier ID

FlashTier0
The tier of high performance flash drives.

FlashTier1
The first tier of high capacity flash drives.

FlashTier2
The second tier of high capacity flash drives.

ENT Enterprise tier consists of drives with speeds of 10K RPM, 15K RPM, or a combination of drives of both speeds.

NL Nearline (NL) tier consists of high volume SATA or SAS Nearline disk drives.

Unknown
Assigning action was specified, but the target tier is unknown.

Example

Specifying the -pathgrp parameter

If you specify the **-pathgrp** parameter and no path groups exist for this volume, the following message diaplays:

```
CMUC00234I lsckdvol: No Path Groups found.
```

If you specify the **-pathgrp** parameter and path groups exist for this volume, a path group status table appends to the resulting output.

```
dscli> showckdvol –pathgrp efff
```

The resulting output

```
Name              efff
ID                EFFF
accstate          Online
datastate         Normal
configstate       Normal
...
migrating         0
perfgrp           PG31
migratingfrom     -
resgrp            RG62
==========================Path Group status=======================
```

```
GroupID                  State       Reserve    Mode        Sysplex
================================================================
800002B9472827CA78BC17   Ungrouped   Disabled   Single      N/A
880005B9472827CAAD6FBA   Grouped     Disabled   Multi-path  N/A
800009B9472827CAC684B9   Ungrouped   Disabled   Single      PLEXM1
```

Report field definitions (*-pathgrp* parameter is specified)

GroupID
The path group ID. An 11-byte value that is displayed as 22 hexadecimal characters.

Note: The path group ID is supplied by the host and is not interpreted further by the DS8000 system. This process means that the hosts are free to define, or redefine, the meaning of this value with no impact to the DS8000 system. However, some programs such as ICKDSF break drown the ID into distinct fields with the following partial display of the ICKDSF logical path status table.

```
+--------------------------+-
|       PATH GROUP ID      | ...
|------+------+----+--------+-
|      |      |CPU |CPU TIME|
|  ID  |SERIAL|TYPE|  STAMP | ...
+------+------+----+--------+-
|800002|B947  |2827|CA78BC17| ...
+------+------+----+--------+-
|880005|B947  |2827|CAAD6FBA| ...
+------+------+----+--------+-
|800009|B947  |2827|CAC684B9| ...
+------+------+----+--------+-
```

For more information, see the *ICKDSF User's Guide and Reference* . Go to http://129.33.205.81/jct03001c/systems/z/os/zos/library/bkserv/v2r1pdf/#ICK to access the ICKDSF PDF.

State The grouped state of this path group. Valid state values are "Grouped" or "Ungrouped".

Reserve
The reserved state of this path group. Valid state values are "Enabled" or "Disabled".

Mode The path mode for this path group. Valid mode values are "Single" or "Multi-path".

Sysplex
The z/OS sysplex name. If the name is not set or available, N/A is displayed.

Logical subsystem specific commands

Commands are referenced for tasks that are associated with Open Systems logical subsystems.

The following logical subsystem specific commands are available:

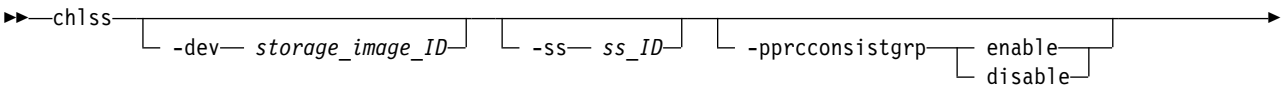
chlss Modifies one or more logical subsystems.

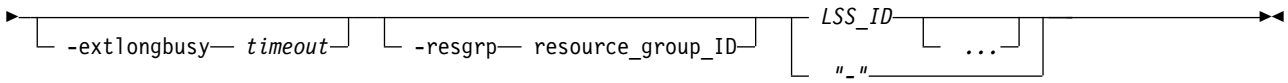
lslss Generates a report that displays a list of logical subsystems (LSSs) for a storage system and the status information for each logical subsystem in the list.

showlss
Generates a report that displays the detailed properties of a specified LSS.

chlss

The **chlss** command modifies a logical subsystem.





Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for all logical subsystems, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-ss *ss_ID*
(Optional). Specifies the subsystem ID of the logical subsystem. The subsystem ID that you specify replaces the existing subsystem ID in the initial target LSS ID.

The ID is in the format 0001 - FFFF.

Example: 013F

-pprcconsistgrp enable | disable
(Optional) Enables a volume that is associated with a logical subsystem to become suspended and enter an extended long busy state if it has not received a notification that a consistency group has been created. Otherwise, the volumes associated with the LSS do not go to a long-busy state.

Note: A consistency group volume pair that is suspended will not go into an extended long busy state unless the primary volume is in a full duplex state before the suspension.

-extlongbusy *timeout*
(Optional) Specifies the time in seconds that an LSS consistency group volume stays long busy after reporting an error that causes a remote mirror and copy suspension if a consistency group has not been created. The default value for new LSSs is 60 seconds (1 minute).

-resgrp *resource_group_ID*
(Optional) Specifies the resource group that the LSSs are assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

LSS_ID ... | **-**
(Required) Specifies one or more LSSs to be modified by this command. An LSS ID is two hexadecimal characters 00 - FE.

To specify a range of LSS IDs, separate the IDs with a hyphen.

You must separate multiple LSS IDs or ranges of LSS IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example: 00-03 08

Example

Invoking the chlss command

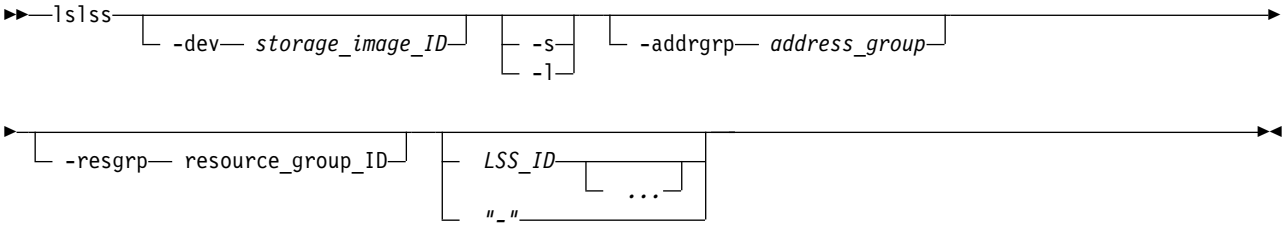
```
dscli> chlss -dev IBM.2107-75FA120 -extlongbusy 120 06 0F
```

The resulting output

```
LSS 06 successfully modified.
LSS 0F successfully modified.
```

lslss

The **lslss** command displays a list of logical subsystems (LSSs) for a storage image and status information for each logical subsystem in the list.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter will temporarily override any defined value for *devid* for the current command.

-s
(Optional) Displays LSS IDs. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output plus the SSID and resource group ID. You cannot use the **-l** and the **-s** parameters together.

-addrgrp *address_group*
(Optional) Displays only LSSs that belong to the specified address group. An address group is a single hexadecimal character (0 - F).

-resgrp *resource_group_ID*
(Optional) Displays only the LSSs that are assigned to the specified resource group ID. The resource group ID begins with the letters *RG* and ends with a decimal number.

LSS_ID ... | **-**
(Optional) Specifies the logical subsystem IDs. An LSS ID is two hexadecimal characters 00 - FE.

To specify a range of logical subsystem IDs, separate the logical subsystem IDs with a hyphen.

You must separate multiple logical subsystem IDs or ranges of logical subsystem IDs with a blank space between each ID or range of IDs.

Example: 00-03 08

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lslss** command using the **-l** parameter.

Invoking the lslss command

```
dscli> lslss -dev IBM.2107-75FA120 -l
```

The resulting output

ID	Group	Addrgrp	Stgtype	Confgvols	Subsys	resgrp
IBM.2107-75FA120/10	0	1	fb	256	FF10	RG0
IBM.2107-75FA120/11	1	1	fb	256	FF11	RG0
IBM.2107-75FA120/12	0	1	fb	256	FF12	RG0

Report field definitions

ID*
Indicates the unique identifier that is assigned to this logical subsystem object. The identifier includes the storage image ID and a 2-digit hexadecimal number for the LSS ID. The LSS ID can be in the range of 00 - FE.

Group
Indicates the server that is managing the logical subsystem group. The server is identified as either 0 or 1.

Addrgrp
Indicates the address group object of which the logical subsystem is a member.

Stgtype
Indicates the type of storage volumes that are contained by this logical subsystem. The displayed value is either fb (fixed block) or ckd (count key data).

Confgvols
Indicates the number of volumes currently assigned to this logical subsystem.

Subsys+
Indicates the user assigned, or default SSID value.

resgrp+
Indicates the resource group ID that the LSS is assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

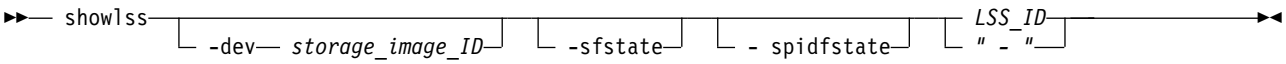
Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

showlss

The **showlss** command displays detailed properties of a logical subsystem (LSS).



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified LSS ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-sfstate
(Optional) Displays a table at the end of the command output that contains the name, ID, and soft fence state for each volume in the specified LSS.

-spidfstate
(Optional) Displays a table at the end of the command output that contains the name, ID, and SPID (Set Path Group ID) fence state for each volume in the specified LCU.

LSS_ID | **-**
(Required) Displays the properties for the specified logical subsystem. This parameter accepts a fully qualified LSS ID, which consists of the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is two hexadecimal digits in the range 00 - FE.

The following example is a fully qualified LSS ID: IBM.2107-75FA120/10

The following example is a shortened version of the LSS ID when the **-dev** parameter is specified:

```
dscli> showlss -dev IBM.2107-75FA120 10
```

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **showlss** command.

Invoking the showlss command to show default information

```
dscli> showlss IBM.2107-75FA120/10
```

The resulting output

ID	Group	Addrgrp	Stgtype	Confgvols
2107-75FA120/10	0	1	fb	256

Subsys	Pprcconsistgrp	Xtndlbztimout	resgrp
FF10	Disabled	60 secs	RG0

Report field definitions

ID Indicates the unique identifier that is assigned to this logical subsystem. It includes the storage image ID and the 2 digit LSS ID 00 - FE.

Group Indicates the server that manages the logical subsystem. The displayed values are 0 or 1.

Addrgrp
Indicates the address group object that the logical subsystem is a member of.

Stgtype
Indicates the type of storage volumes contained by this logical subsystem. The value displayed is fb (fixed block) or ckd (count key data)

Confgvols
Indicates the number of volumes that are assigned to this logical subsystem.

Subsys
Indicates the user assigned, or default SSID value.

Pprcconsistgrp
Indicates the assigned PPRC consistency group state setting. If you do not designate enabled, the default value of disabled is assigned.

Xtndlbztimout
Indicates the assigned extended long busy timeout value. LSSs are created with an extended long busy timeout value of 60 seconds.

resgrp Indicates the resource group ID that the LSS is assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

Example 2

If you specify the **-sfstate** parameter, the output includes the soft fence state table.

```
dscli> showlss -sfstate 14
Date/Time: September 22, 2016 3:18:07 PM MST IBM.
DS: IBM.2107-75LH321
ID              14
Group           0
addrgrp         1
stgtype         fb
confgvols       16
subsys          0xFF14
pprcconsistgrp  Disabled
xtndlbztimout   60 secs
resgrp          RG0
============Soft Fence State============
Name         ID   sfstate
========================
RegrFBVol1 1400 Disabled
RegrFBVol1 1401 Disabled
RegrFBVol1 1402 Disabled
RegrFBVol1 1403 Disabled
RegrFBVol1 1404 Disabled
RegrFBVol1 1405 Disabled
RegrFBVol1 1406 Disabled
RegrFBVol1 1407 Disabled
RegrFBVol1 1408 Disabled
RegrFBVol1 1409 Disabled
RegrFBVol1 140A Enabled
RegrFBVol1 140B Enabled
RegrFBVol1 140C Enabled
RegrFBVol1 140D Enabled
RegrFBVol1 140E Enabled
RegrFBVol1 140F Enabled
dscli>
```

Report field definitions (-*sfstate* parameter is specified)

Name The user-assigned nickname for this volume object.

ID The unique identifier that is assigned to this volume object. A volume ID is four hexadecimal characters (0x0000 - 0xFEFF).

sfstate The soft fence state. The following values are possible.

Enabled
The host has set this volume to the soft fence state.

Disabled
The host has not set this volume to the soft fence state.

N/A The host cannot set this volume to the soft fence state. For example, an alias volume.

Report field definitions (`-spidfstate` parameter is specified)

Name The user-assigned nickname for this volume object.

ID The unique identifier that is assigned to this volume object. A volume ID is four hexadecimal characters (0x0000 - 0xFEFF).

spidfstate
The soft fence state. The following values are possible.

Enabled
The volume is set to the SPID fence state by the host.

Disabled
The volume is not set to the SPID fence state by the host.

N/A The volume is not capable of being set into the SPID fence state by the host. For example, an alias volume.

Fixed block logical volume specific commands

Commands are referenced for tasks that are associated with Open Systems fixed block logical volumes.

The following fixed block logical volume specific commands are available:

`chfbvol`
Modifies the name or data type of a fixed block volume.

`initfbvol`
Releases extents from a space-efficient logical volume, and removes (erases) data from standard volumes.

`lsfbvol`
Generates a report that displays a list of fixed block volumes in a storage system and the status information for each volume in the list.

`managefbvol`
Initiates a change on fixed block (FB) volumes by executing a process.

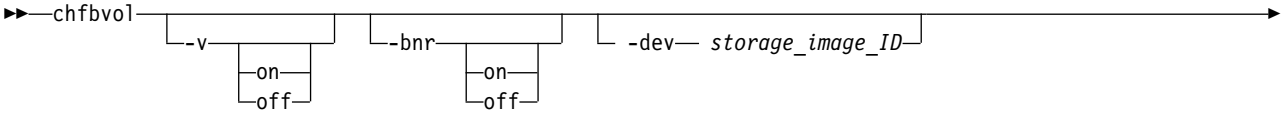
`mkfbvol`
Creates open systems fixed block volumes in a storage system.

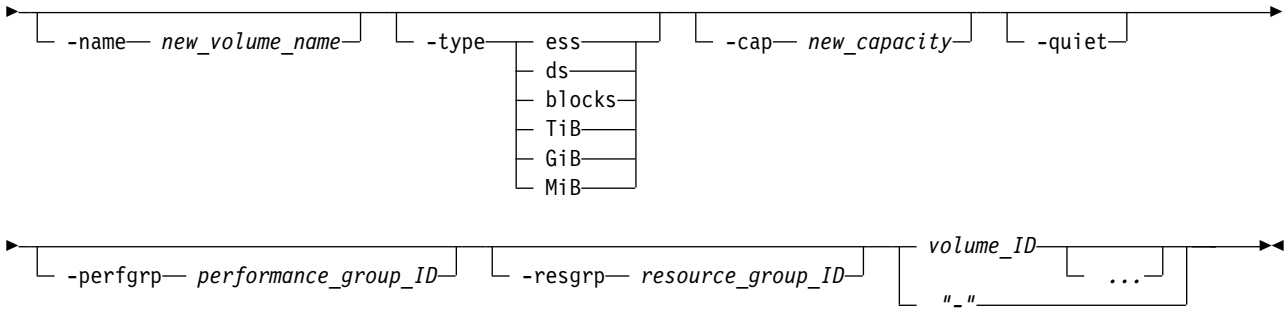
`rmfbvol`
Deletes one or more specified fixed block volumes from a storage system.

`showfbvol`
Generates two types of reports. The first report displays the detailed properties for a specified fixed block volume. The second report displays the performance metrics for a specified fixed block volume.

chfbvol

The **`chfbvol`** command is used to change the name or capacity of a fixed block volume.





Parameters

-v on | off
(Optional) Turns verbose mode on or off. The default is off. You can specify `-v on` or `-v off` to turn verbose mode on or off; but if you specify only `-v`, you get an error.

-bnr on | off
(Optional) Turns the banner on or off. The default is on. You can specify `-bnr on` or `-bnr off` to turn the banner on or off; but if you specify only `-bnr`, you get an error.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of the manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-name *new_volume_name*
(Optional) Specifies the nickname for this volume. A nickname cannot exceed 16 characters.

-type ess | ds | blocks
(Optional) Specifies the unit type of capacity that is specified by the **-cap** parameter.

ess Specifies that the unit is 10^9 bytes.

ds Specifies that the unit is 2^30 bytes.

blocks Specifies that the unit is 512 blocks.

TiB Specifies that the unit is tebibytes (TiB) 2^40 bytes.

GiB Specifies that the unit is gibibytes (GiB) 2^30 bytes.

MiB Specifies that the unit is mebibytes (MiB) 2^20 bytes.

Note: If the **-type** parameter is not specified, the LUN is created as type *ds*.

-cap *new_capacity*
(Optional) Specifies the storage size that you want to allocate to the specified volume. Check your operating system documentation to ensure that volume expansion is supported before you proceed with the expansion.

- If the **-type** parameter is omitted or the **-type** *ds* parameter is specified, the new_capacity value is the volume size in gibibytes (GiB), where 1 gibibyte (GiB) = 1 073 741 824 (2^30 bytes).
- If the **-type** *ess* parameter is specified, the new_capacity value is the volume size in gigabytes (GB), to the nearest 1/10 GB (format xxxx.x), where one GB = 1 000 000 000 (10^9 bytes).
- If the **-type** *blocks* parameter is specified, the new_capacity value is the volume size in 512 byte blocks.

Notes:

1. The maximum volume size varies and depends on DS8000 model and type.
2. If you attempt to reduce the volume size, the command fails.

-quiet
(Optional) Turns off the confirmation prompt for this command.

-perfgrp *performance_group_ID*
(Optional) Specifies the performance group ID that the volumes are assigned to. The performance group ID begins with the letters *PG*. The default is *PG0*.

-resgrp *resource_group_ID*
(Optional) Specifies the resource group that the volumes are assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

volume_ID ... | -
(Required) Specifies one or more volume IDs to be modified. The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ* where:

X Specifies the address group, 0 - F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

You must fully qualify the volume ID with manufacturer, machine type, and serial number if you do not use the **-dev** parameter.

To specify a range of volume IDs, separate the volume IDs with a hyphen.

You must separate multiple volume IDs or ranges of volume IDs with a blank space between each ID or range of IDs.

Example: 0100-010F 0180-018F 0120

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive command mode.

Example 1

Invoking the chfbvol command

```
dscli> chfbvol -dev IBM.2107-75FA120 -cap 5 0100 0101
```

The resulting output

```
FB volume 0100 successfully modified.
FB volume 0101 successfully modified.
```

Example 2

Invoking the chfbvol command

```
dscli> chfbvol -cap 2 -quiet  0510
```

The resulting output

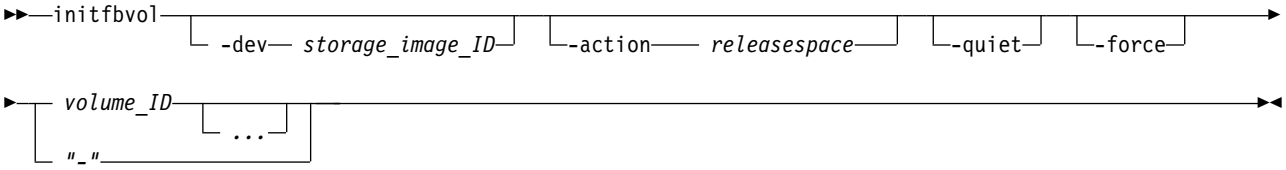
```
| CMUC00332W chfbvol:
| Some host operating systems do not support changing the volume size.
| Data can be at risk if the host does not support this action.
| Are you sure that you want to resize the volume? [y/n]: y
  CMUC00026I chfbvol: FB volume 0510 successfully modified.
```

initfbvol

The **initfbvol** command initializes a logical volume and releases extents from a space-efficient logical volume.

For space efficient logical volumes, this command is used to release space. For TSE volumes, it releases tracks in the repository, reducing the repository allocated space. For ESE volumes, it releases extents in the extent pool being used, reducing the allocated extents.

For example, if a space-efficient logical volume has data that is stored on it that is no longer needed, use the **initfbvol** command to free the extents/tracks that were assigned to this logical volume. This allows the extents/tracks to be reused by other space-efficient logical volumes.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a value for manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-action releasespace
(Optional) Specifies that you want the system to release the repository space held by the designated space-efficient volume back to the repository. (The repository is the physical extents that provision the virtual extents for virtual space volumes.)

-quiet
(Optional) Turns off the modification confirmation prompt for the **-action** parameter.

-force
(Optional) Forces the action.

Note: For release 8.3 or later, if **-force** is not specified, the request fails when the volume is in a volume group or in any kind of Copy Services relationship.

volume_ID **-** | **...**
(Required) Specifies the volume ID that you want the system to release the repository space from. The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *LLVV*, where:

LL Specifies the logical control unit number, 00 - FE.

VV Specifies the volume number, 00 - FF.

You must fully qualify the volume ID with manufacturer, machine type, and serial number if you do not use the **-dev** parameter.

To specify a range of volume IDs, separate the volume IDs with a dash.

You must separate multiple volume IDs or ranges of volume IDs with a blank space between each ID or range of IDs.

If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive command mode. The ellipsis (...) indicates that, optionally, you can specify multiple values.

Example

Invoking the initfbvol command

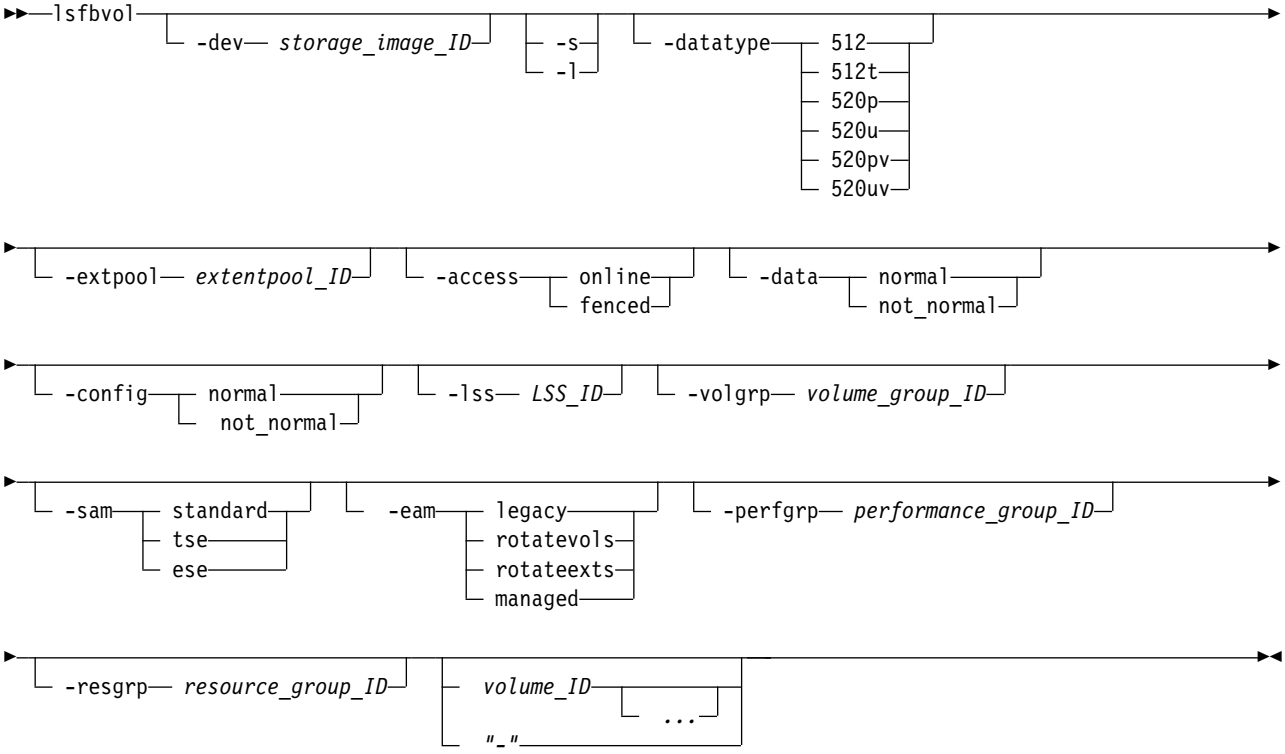
```
dscli> initfbvol -dev IBM.2107-75FA120 -action releasespace 0101
```

The resulting output

```
CMUC00338W initfbvol: Are you sure that you want to submit the command
 releasespace for the FB volume 0101?[Y/N]:y
CMUC00341I initfbvol:: 0101: The command releasespace has
completed successfully.
```

lsfbvol

The **lsfbvol** command displays a list of fixed block (FB) volumes in a storage image and status information, including Safeguarded Copy information for each volume in the list.



Parameters

Note: For a storage system that is heavily configured, specify the **-lss** or the **-volgrp** parameter as part of your command.

-dev *storage_image_ID*
(Optional) Displays the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs. It is also required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-s
(Optional) Displays only the volume IDs. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays default output plus extra attributes that are identified as long output in the Report field definitions list. You cannot use the **-l** and the **-s** parameters together.

-datatype 512 | 512t | 520p | 520u | 520pv | 520uv
(Optional) Displays volumes of the specified volume data type. Standard 2107/1750 volume (512), T10-DIF algorithm protection (512t), System i protected (520p), System i unprotected (520u), IBM i protected variable size (520pv), IBM i unprotected variable size (520uv).

-extpool *extentpool_ID*
(Optional) Displays volumes that are sourced from the specified extent pool. An extent pool ID is a four-digit decimal number with no leading zeros, prefixed with the letter P.

-access online | fenced
(Optional) Displays volumes with the specified access state.

-data normal | not_normal
(Optional) Displays volumes with the specified data state.

-config normal | not_normal
(Optional) Displays volumes with the specified configuration.

-lss *LSS_ID*
(Optional) Displays volumes with IDs that contain the specified logical subsystem ID. Each logical subsystem can contain up to 256 volumes. A logical subsystem ID is two hexadecimal characters 00 - FE.

-volgrp *volume_group_ID*
(Optional) Displays volumes that are assigned to the specified volume group ID. A volume group ID is a four-digit decimal number, with no leading zeros, prefixed by the letter *V*. For example, V123.

-sam standard | tse | ese
(Optional) Specifies the storage allocation method as follows.

standard
Specifies that you want the system to fully allocate the volume with real extents when it creates the volumes.

tse Specifies that you want the system to create track space-efficient volumes. After creation, these space-efficient volumes contain a set of virtual extents that are associated with the space-efficient storage in the same extent pool. The physical space for a given logical track on a track space-efficient logical volume is dynamically allocated and deallocated from the repository in the space-efficient storage.

Note: Track space-efficient (TSE) repositories are required to create TSE volumes. However, starting with DS8000 Release 8, TSE repositories can no longer be created.

ese Specifies that an extent space efficient logical volume is provisioned with a set of virtual extents that are associated with the space efficient storage in the same extent pool. Physical space for an extent space efficient logical volume is dynamically allocated and deallocated from the extent pool.

-eam legacy | rotatevols | rotateexts | managed
(Optional) Specifies that you want the system to display only volumes that meet the criteria of the designated extent allocation method as follows:

legacy Specifies that the volumes that were created before the current algorithms were implemented.

rotateexts
Specifies that the extents for each new logical volume are allocated across all available ranks, and is also known as storage-pool striping. This value is the default.

rotatevols
Specifies that the extents for each new logical volume are allocated from each successive rank. The extents for a particular volume will be allocated from one rank, while the extents for the next volume will be allocated from the next successive rank, and so on.

managed
Specifies that the extents are currently managed by Easy Tier, and the extents for any new volumes are initially allocated across all available ranks in the lowest tier of storage.

-perfgrp *performance_group_ID*
(Optional) Displays only the volumes that belong to the specified performance group. The performance group ID begins with the letters *PG*.

-resgrp *resource_group_ID*
(Optional) Displays only the volumes that are assigned to the specified resource group ID. The resource group ID begins with the letters *RG* and ends with a decimal number.

volume_ID ... | –
(Optional) Displays volumes with the specified IDs. The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ* where:

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

X Specifies the address group, 0–F.

To specify a range of volume IDs, separate the volume IDs with a hyphen.

You must separate multiple volume IDs or ranges of volume IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 0100-010F 0180-018F 0120

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lsfbvol** command using the **-l** parameter.

Invoking the lsfbvol command

```
dscli> lsfbvol -dev IBM.2107-75FA120 -l -volgrp V2
```

The resulting output

Name	ID	accstate	datastate	config state	device MTM	data type
My_ volume_ 0001	0100	Online	Normal	Normal	2107-900	FB 512

Name	ID	accstate	datastate	config state	device MTM	data type
My_ volume_ 0002	0102	Online	Normal	Normal	2107-A07	FB 520P
My_ _volume 0003	0103	Online	Normal	Normal	2107-900	FB 512
My_ _volume 0004	0104	Online	Normal	Normal	2107-099	FB 520UV
My_ _volume 0007	0105	Online	Normal	Normal	2107-050	FB 520PV

extpool	sam	captype	cap (2^30B)	cap (10^9B)	cap (blocks)	volgrp
P21	standard	DS	64.0	-	134217728	V2
P31	standard	iSeries	128.0	68.7	268435456	V2
P21	standard	ESS	-	35.0	68359424	-
P0	standard	iSeries	5.0	5.4	10485760	-
P0	standard	iSeries	5.0	5.4	10485760	-

eam	reqcap (blocks)	perfgrp	resgrp
legacy	2097152	PG1	RG0
rotateexts	2097152	PG1	RG0
legacy	2097152	PG1	RG0
rotateexts	10485760	PG0	RG0
rotateexts	10485760	PG0	RG0

GUID	safeguarded ap (2^30B)	safeguardedloc	usedsafeguardedcap (2^30B)	safeguarded
60050763CXFFC03 D0000000000000100	16.0	mypool_1	1.0	yes
60050763CXFFC03 D0000000000000101	16.0	mypool_1	1.0	yes
60050763CXFFC03 D0000000000000102	16.0	mypool_1	1.0	yes
60050763CXFFC03 D0000000000000103	16.0	mypool_1	1.0	yes

Report field definitions

Name
Indicates the nickname that you assigned for the specified volume object.

ID Indicates the unique identifier that is assigned to this volume object.

Accstate
One of the following access states are displayed: Online or Fenced.

Online
The logical volume is accessible to a host.

Fenced
The logical volume is in the volume fenced state and is not accessible to the host.

Datastate
One of the following data states are displayed:

Normal
Indicates that none of the other data states apply. The access state is Online.

Read only
Indicates that the logical volume is read only because one or more extents on the logical volume are on a rank in the read only data state. The access state is Online.

Inaccessible
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the inaccessible data state. The access state is fenced.

Virtual space fault
Indicates that the logical volume has a storage allocation method of track space-efficient. Not enough space was available to convert a virtual logical track to a real logical track. The access state is Online.

Indeterminate data loss
Indicates that the following data states do not apply and that one of the following conditions occurred:

Data states that do not apply:

- Rank failed
- Rank repairing
- Rank repaired
- Global inaccessible
- Global lost data

Conditions - one of the conditions occurred:

- Committed write data was lost before it was destaged and the track identifiers that are associated with the data are unknown.
- Data was lost that indicates that extents on the logical volume were active FlashCopy targets.

The access state is fenced.

Rank failed
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the Failed data state. The access state is Fenced. This data state transitions to the Rank repairing state if the rank transitions to the Rank repairing state through use of the repair array function.

Rank Repairing
Indicates that one or more extents that are associated with the logical volume are on ranks in the repairing data state. The access state is fenced.

Rank Repaired
Indicates that one or more extents that are associated with the logical volume are on ranks that were in the repairing state, but are not in the repairing state now. The access state is fenced.

Global inaccessible
Indicates that the global metadata that is associated with the logical volume configuration is inaccessible. Some of the data that is associated with the logical volume might be inaccurate. The access state is fenced.

Global lost data
Indicates that global metadata that is associated with the logical volume configuration was lost. As a result, some of the data that is associated with the logical volume might be inaccurate. The access state is fenced.

NVS data inaccessible
Indicates that active NVS data is inaccessible for one or more logical volumes of an LSS group. The logical volumes in the LSS group cannot be made accessible. The access state is fenced.

Configstate
One of the following configuration states are displayed:

Normal
Indicates that no logical volume configuration operations are in progress, and the volume is not being deconfigured, merged, or migrated.

Configuring
Indicates that the logical volume is in the process of being configured for the first time.

Reconfiguring
Indicates that the logical volume is allocating or deallocating extents due to a modification of the requested capacity attribute after initial creation.

Deconfiguring
Indicates that the logical volume is being deleted.

Configuration error
Indicates that the initial configuration did not complete successfully. This state reflects an internal error condition and not an error in the request to create the volume. If you have a volume in this state, use the **rmfbvol** command to delete each volume that is listed with the configuration state of "configuration error".

Merging
Indicates that the volume is merging. For example, merging from one extent pool to a different extent pool.

Migrating
Indicates that the volume is migrating, or waiting to be migrated.

Migration Cancelled
Indicates that the volume was migrating and then the ‘migcancel' action of the **manageckdvol** command was issued. Some of the extents were waiting to be migrated in the source pool while other extents already migrated to the target pool. Migration stopped, and cannot be resumed. If you have a volume in this state, try to migrate it again to the original source or target extent pool.

Migration Paused
Indicates that the volume was migrating and then the ‘migpause' action of the **manageckdvol** command was issued. Migration stopped, but can be resumed.

Migration Error
Indicates that the volume migration process failed to complete successfully. This state reflects an internal error condition and not an error in the request of the user to migrate a volume. If you have a volume in this state, try to migrate it again to the original source or target extent pool.

Reconfiguration error
Indicates that the reconfiguration request did not complete successfully.

Deconfiguration error
Indicates that a request to delete a volume did not complete successfully. This state reflects an internal error condition and not an error in the request to remove the volume. To correct this state, you must reissue the **rmfbvol** command for the designated volume.

Transposition Error
Indicates that the volume is in an extent pool that was unsuccessfully merged. This state reflects an internal error condition. **Corrective action** Use the **chextpool** command with the **-merge** parameter again to redrive the merge extent pool and to correct this state.

deviceMTM
Indicates the volume device type and model. The volume MTM (machine type, model) is determined by the fixed block volume data type and the provisioned capacity (in GB). The machine type is either 2107 or 1750; however, the MTM can be any one of the following, depending on your system:

2107-900
Indicates a standard 2107 volume.

1750-500
Indicates a standard 1750 volume.

xxxx-A0*x*
Indicates that the *xxxx* is a 2107 or 1750. The value A0 indicates a System i protected volume (for example, 2107-A01 or 1750-A07).

xxxx-A8*x*
Indicates that the *xxxx* is 2107 or 1750. The value A8 indicates a System i unprotected volume (for example, 2107-A81 or 1750-A87).

2107-050
Indicates that the machine type is 2017. The value 050 indicates that the volume is a System i unprotected variable size volume.

2107-099
Indicates that the machine type is 2107. The value 099 indicates that the volume is a System i protected variable size volume.

Datatype
Indicates the volume data type setting. One of the following values is displayed:

- FB 512
- FB 512T
- FB 520P
- FB 520U
- FB 520PV
- FB 520UV

Extpool
Indicates the extent pool ID. Volume extents are allocated from this extent pool ID.

Note: Volumes that belong to an encrypted extent pool are encrypted. You can see the key group of an extent pool by using the **lsextpool -l**, or **showextpool** commands.

SAM
Indicates the storage allocation method. The following values are displayed:

standard
Indicates that the system fully allocated the volume with real extents at volume creation time.

tse
Indicates that a track space-efficient logical volume contains a set of virtual extents that are associated with the space-efficient storage in the same extent pool. Physical space for a given logical track on a track space-efficient logical volume is dynamically allocated and deallocated from the repository in the space-efficient storage.

ese
Indicates that an extent space efficient logical volume is provisioned with a set of virtual extents

that are associated with the space efficient storage in the same extent pool. Physical space for an extent space efficient logical volume is dynamically allocated and deallocated from the extent pool.

Captype
Indicates the capacity unit type that is used at volume creation. One of the following values is displayed:

ESS
The capacity unit is decimal gigabytes (GB).

DS The capacity unit is gibibytes (GiB).

DS/ESS
The capacity unit is gibibytes (GiB) or decimal gigabytes (GB).

Blocks
The capacity unit is 512B.

iSeries
The capacity unit was not specified at volume creation. This fixed block volume was created only for iSeries.

Cap (2^30B)
Indicates the size of the volume that is available for host system access in gibibytes (GiB).

Note: " – " is displayed if the capacity unit type of the volume is ESS (captype=ESS)

Cap (10^9B)
Indicates the size of the volume that is available for host system access in decimal gigabyte (GB) units.

Note: " – " is displayed if the capacity unit type of the volume is DS (captype=DS)

Cap (blocks)
Indicates the quantity of volume logical blocks that are available for host system access.

Volgrp
Indicates the volume groups (excluding default volume groups) that a volume belongs to.

Multiple volume groups that are associated with the volume are separated by a comma.

A " – " is displayed if no volume groups exists that are associated with the volume.

`Unknown` displays if information about the volume groups is not available.

EAM
Indicates the extent allocation method that is used if the volume is migrated or expanded.

legacy
Indicates that the volume was created before the use of the current algorithm.

rotateexts
Indicates that the extents for each new logical volume are allocated across all available ranks, and is also known as storage-pool striping. This value is the default.

rotatevols
Indicates that the extents for each new logical volume are allocated from each successive rank. This means that the extents for a particular volume will be allocated from one rank, while the extents for the next volume will be allocated from the next successive rank, and so on.

managed
Indicates that the extents are currently managed by Easy Tier, and the extents for any new volumes are initially allocated across all available ranks in the lowest tier of storage.

" _ "
A dash " - " value is displayed if the extent allocation method does not apply , for example if the volume is a track space efficient (TSE) volume.

Reqcap (blocks)
Indicates the requested quantity of volume logical blocks (for example, 3339).

perfgrp
Indicates the performance group ID that the volume is assigned to. The performance group ID begins with the letters *PG* and ends with a decimal number.

resgrp
Indicates the resource group ID that the volume is assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

GUID
Specifies the GUID (Globally Unique Identifier): The NAA (Network Addressing Authority) unique value that identifies a volume object.

safeguardedcap(2^30B)
Specifies the amount of safeguarded backup capacity that is reserved for safeguarded copies.

safeguardedloc
Specifies the pool where the safeguarded copies are written to.

usedsafeguardedcap(2^30B)
Specifies the amount of physical pool capacity that is used by the safeguarded copies.

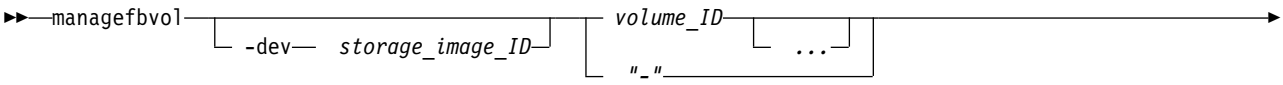
safeguarded
Specifies whether Safeguarded Copy is enabled for the volume.

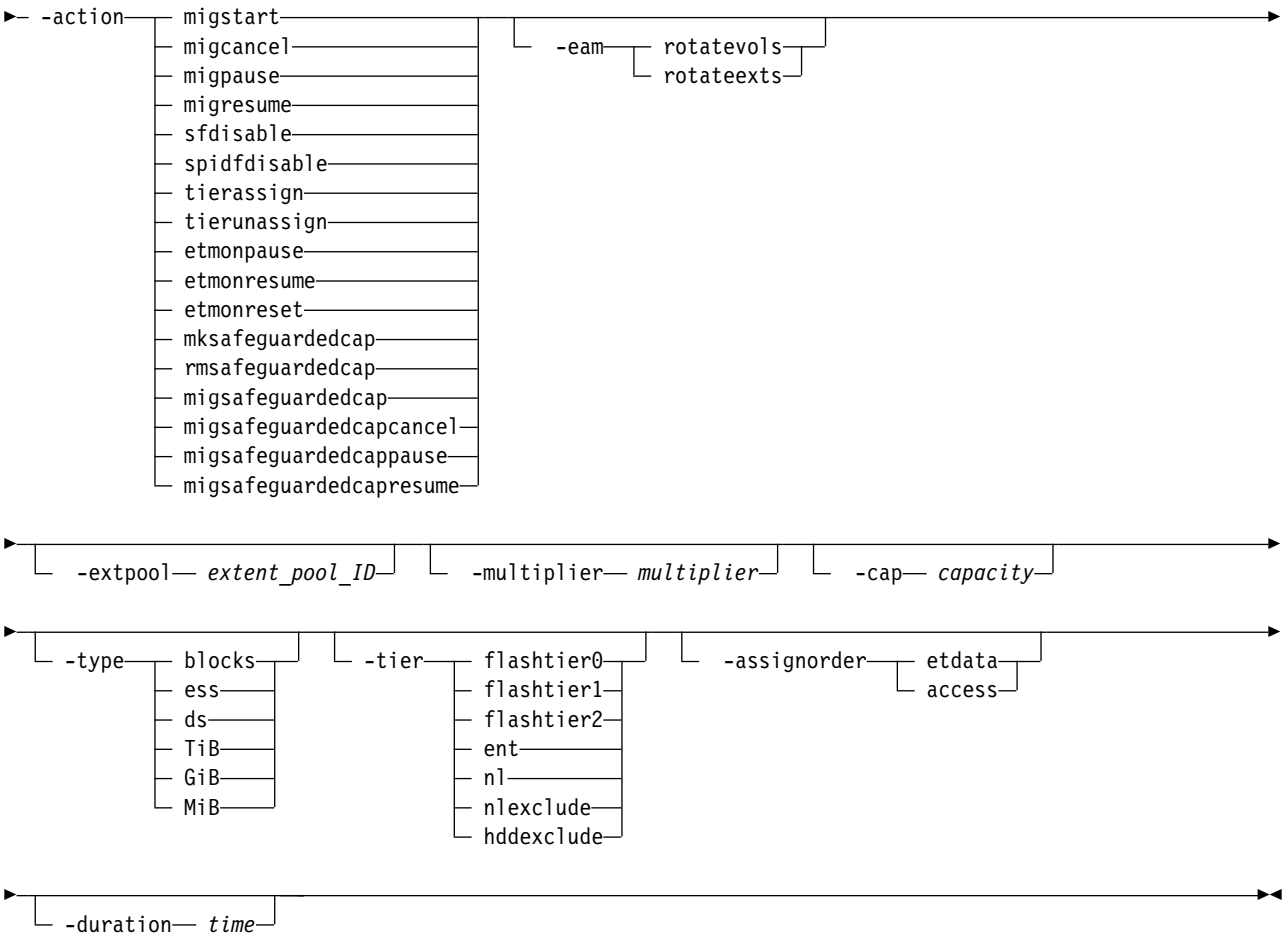
managefbvol

The **managefbvol** command initiates a change on fixed block (FB) volumes when the command processes.

Use this command to define safeguarded backup capacity for FB volumes for which you want to enable for Safeguarded Copy.

Use Copy Services Manager to define a safeguarded copy schedule to create backups on a regular basis, such as hourly or daily.





Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

volume_ID ... | -
(Required) Specifies one or more volume IDs to be modified. The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ* where:

X Specifies the address group, 0 - F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

You must fully qualify the volume ID with manufacturer, machine type, and serial number if you do not use the **-dev** parameter.

To specify a range of volume IDs, separate the volume IDs with a hyphen. Separate multiple volume IDs or ranges of volume IDs with a blank space between each ID or range of IDs.

Example: 0100-010F 0180-018F 0120

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive command mode.

-action
migstart|migcancel|migpause|migresume|sfdisable|spidfdisable|tierassign|tierunassign|etmonpause|etmonresume|etmonreset|mksafeguardedcap|rmsafeguardedcap|migsafeguardedcap|migsafeguardedcapcancel|migsafeguardedcappause|migsafeguardedcapresume

(Required) Specifies that one of the following actions are to be performed:

migstart
Initiates volume migration on the specified volumes that are in the "normal" or "canceled" state. Volumes are placed into the migrating state. Volumes that are in the cancelled state must have the original source or the destination extent pool as the value of the **-extpool** parameter.

migcancel
Cancels volume migration on the specified volumes that are in the migrating state. Volumes that have not yet started migration are put in the "normal" state, and volumes that are in the middle of migration are put in the "canceled" state.

migpause
Pauses volume migration on the specified volumes that are in the migrating state. Volumes that have not yet started migration or that are in the middle of migration are put in the "paused" state.

migresume
Resumes volume migration on the specified volumes that are in the "paused" state.

sfdisable
Sends a soft fence reset command to each specified volume. This action cannot be used with any other parameter.

spidfdisable
Sends a set path group ID (SPID) fence reset command to each specified volume. This action cannot be used with any other parameters.

tierassign
Initiates the assigning volume action on the specified volumes to the specified tier. The **-tier** option is required with this action.

Note: By assigning a volume to a tier, a process that migrates the data in the volume to the specified tier begins. You can check the progress of the migration by using the **showfbvol** or **showckdvol** command with the **-tier** parameter. However, there is a maximum capacity that can be assigned to each tier in a pool. If this maximum capacity is reached, a new assign command to that tier is rejected due to insufficient capacity. You can check how much total capacity has been assigned to a specific tier in a pool by using the **showextpool** command with the **-tier** parameter. If, due to a configuration change, the assign operation reaches the maximum capacity on the target tier before all volumes are assigned, the remaining volumes enter the `Assign Pending Hardware` state. The migration of those volumes to the target tier are paused.

tierunassign
Initiates the unassigning volume action on the specified volumes.

etmonpause
Specifies that Easy Tier monitoring of this volume will be paused. During the pause, all Easy Tier storage migrations are unaffected, but no new migration plans will be formed.

etmonresume
Specifies that Easy Tier monitoring of this volume will be resumed. All Easy Tier storage migrations are unaffected.

etmonreset
Specifies that all Easy Tier monitoring data (history), including migration plans are erased. All new plans will be based on new monitoring data.

mksafeguardedcap
Creates the safeguarded backup capacity for one or more of the specified volumes.

rmsafeguardedcap
Deletes the safeguarded backup capacity, including all safeguarded copies, that was defined for one or more of the specified volumes.

migsafeguardedcap
Migrates the safeguarded backup capacity to a new pool for one or more of the specified volumes.

migsafeguardedcapcancel
Cancels migrating the safeguarded capacity to a new pool for one or more of the specified volumes.

migsafeguardedcappause
Pauses migrating the safeguarded backup capacity to a new pool for one or more of the specified volumes.

migsafeguardedcapresume
Resumes migrating the safeguarded backup capacity to a new pool for one or more of the specified volumes.

-eam
(Optional) Specifies the extent allocation method as follows:

rotateexts
Designates that extents that are allocated to a logical volume are successively rotated through the ranks within an extent pool. This parameter is the default value.

rotatevols
Designates that each successive logical volume that is created is allocated on the next available rank in the extent pool.

Note: You can specify the **-eam** parameter only if **-action** migstart is also specified.

-extpool *extent_pool_ID*
(Optional) Changes the extent pool ID of the volume so that the volume migrates to the new extent pool. Accepts either a fully qualified extent pool ID including storage image ID or a shortened version if the -dev parameter is used. The shortened version is a four-digit decimal number with no leading zeroes, prefixed with the letter P.

Note: When the command returns, the volume migration might still be occurring. It is available for I/O and copy services during migration. Its configstate indicates that it is migrating.

-multiplier *multiplier*
(Optional) Specifies the multiplier value. Selecting the correct backup capacity multiplier depends on the provisioned capacity being backed up, the number of copies, the interval over which the copies are kept, and the predicted destage rate of the volume data.

-cap *capacity*
(Optional) Specifies the safeguarded backup capacity for one or more specified volumes. The unit type is set with the parameter **-type**, if the unit type is not set. The default is *GiB*.

-type blocks | ess | ds |TiB | GiB | MiB
(Optional) Specifies the unit type of the safeguarded backup capacity. If the unit type is not set, the default is *GiB*.

blocks Specifies the capacity unit type as blocks. This capacity unit type can be selected only for fixed block storage. One block = 512 Bytes.

ess Specifies that the unit is 10^9 bytes.

ds Specifies that the unit is 2^30 bytes.

TiB Specifies that the unit is tebibytes (TiB) 2^40 bytes.

GiB Specifies that the unit is gibibytes (GiB) 2^30 bytes.

MiB Specifies that the unit is mebibytes (MiB) 2^20 bytes.

-tier flashtier0 | flashtier1 | flashtier2 | ent | nl | nlexclude | hddexclude
(Optional) Specifies which tier the volume is assigned to. This option is required with the **-action tierassign** parameter.

flashtier0
The tier of high performance flash drives.

flashtier1
The first tier of high capacity flash drives.

flashtier2
The second tier of high capacity flash drives.

ent Enterprise tier that consists of drives with speeds of 10K RPM, 15K RPM, or a mixtures of 10K RPM and 15K RPM speeds.

nl Nearline tier consists of high-volume disks that are either SATA or SAS Nearline drives.

NLExclude
Enterprise tiers but not a Nearline tier.

HDDexclude
Any tier except enterprise (ent) or nearline (nl).

-assignorder etdata|access
(Optional) Specifies the order in which the data is migrated. This option is valid only with the **-action tierassign** parameter.

etdata While all data is scheduled to migrate, the migration order is based on the prioritization of the data as specified in the Easy Tier heat map. This value allows the specified volume to be pre-staged onto the specified tier. This is the default value if **-assignorder** is not specified.

access While all data is scheduled to migrate, the data is migrated only when accessed. In other words, data that is never accessed is never migrated to the specified tier.

-duration *time*
(Optional) Specifies the hours of the pause time in ISO 8601 format. For example, **-duration** 24H. The maximum value of the time is a week, which is 168 hours (168H). You can specify this option only with **-action etmigpause** or **etmonpause** parameters.

Note: If you want the duration of the pause to be infinite, you must specify **-duration** 0H. Otherwise, if you do not specify a value with the **-duration** parameter, the default is 168H.

Example

Invoking the managefbvol command

```
dscli> managefbvol -dev IBM.2107-75FA120
-action migstart  -extpool P2 0100
```

The resulting output

```
CMUC00000I manageckdvol:  CKD Volume 0100 action migstart executed successfully.
```

Invoking the managefbvol command

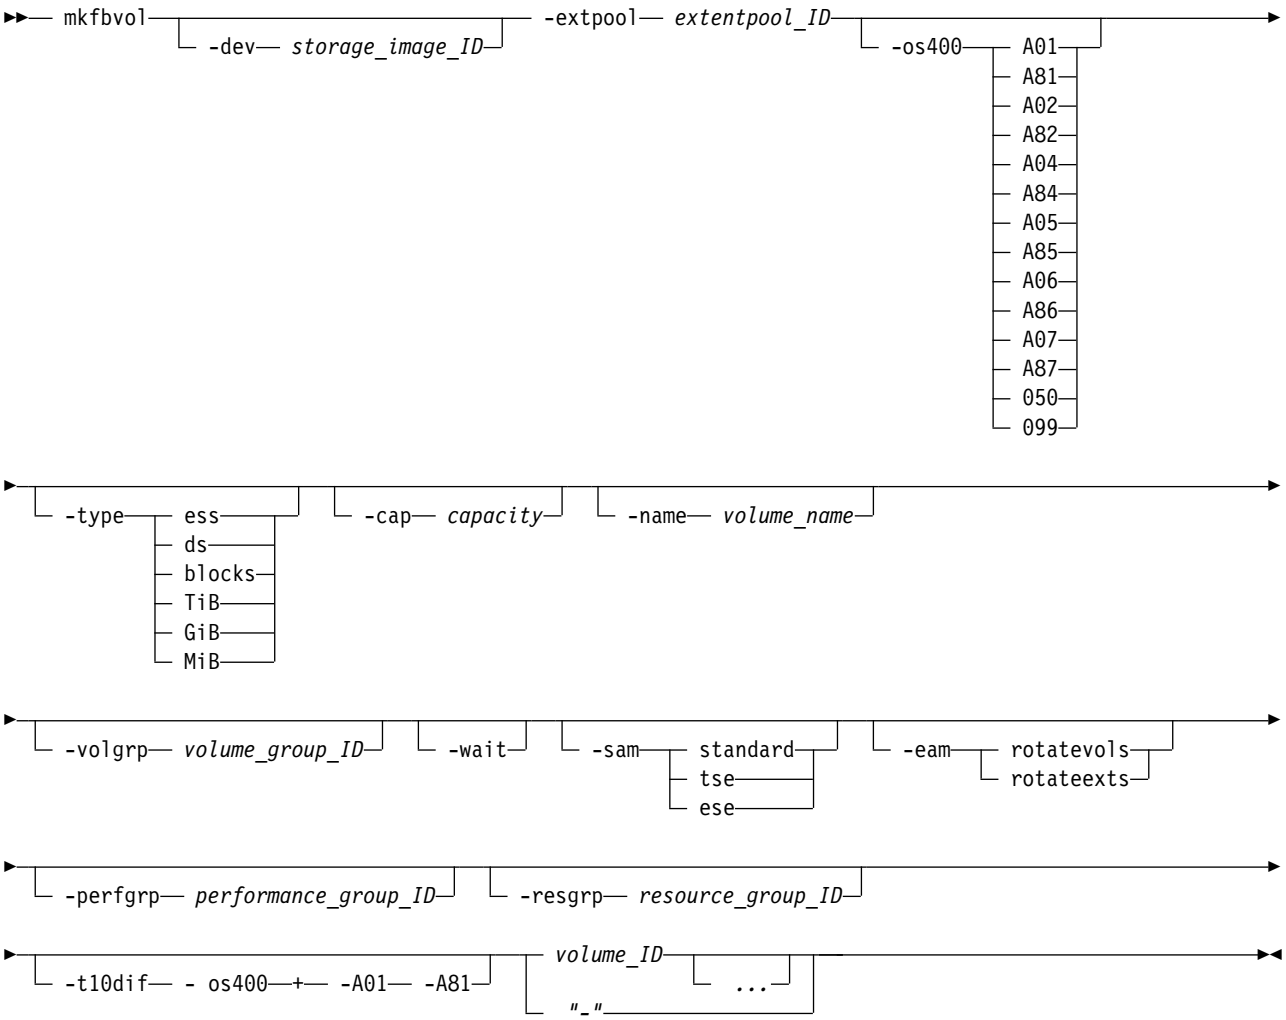
```
dscli> managefbvol -dev IBM.2107-75FA120 -action rmsafeguardedcap 000F
```

The resulting output

```
CMUC00568W managefbvol: Are you sure you want to delete the backup capacity for FB volume 000F? [y/n]: y
CMUC00573I managefbvol: Backup capacity for FB volume 000F is successfully deleted.
```

mkfbvol

The **mkfbvol** command creates open systems fixed block (FB) volumes in a system.



Parameters

Notes:

1. You can create multiple volumes with one command; however, all volumes must have the same capacity, extent pool, and data type.
2. If host attachment volume groups are not created, create temporary volume groups and assign new fixed block volumes to the temporary volume groups. Create the volume groups according to the volume type and capacity characteristics.

3. To use the -sam *tse* parameter, you must have previously created space-efficient storage (using the **mksestg** command) for the extent pool.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes a value for the manufacturer, machine type, and serial number. The **-dev** parameter temporarily overrides any defined value for *devid* for the current command. The storage image ID is required for the following situations:

- If you do not specify a fully qualified extent pool ID
- If you do not set the *devid* variable in your profile or through the **setenv** command
- If the HMC is aware of more than one storage image

-extpool *extentpool_ID*
(Required) Creates the base or alias volumes from data extents that are contained in this extent pool. The extent pool storage type defines the volume storage type. An extent pool ID is a four-digit decimal number with no leading zeros, prefixed with the letter *P*.

-os400 A01 | A81 | A02 | A82 | A04 | A84 | A05 | A85 | A06 | A86 | A07 | A87 | 050 | 099
(Optional) The OS/400 volume options. If this parameter is not specified, the default standard 2107/1750 volume is created.

This parameter, with a specified *Axx* value, is required if the **-cap** parameter is not specified because the *Axx* values also indicate a specific capacity. However, because the *0xx* values do not indicate any capacity, the **-cap** parameter is required with the *0xx* values to specify the capacity of the IBM i volume.

The storage sizes and the data types for this volume are listed in the following table:

Volume	Volume Size	Volume	Volume Size
A01 protected	8.59 GB	A81 unprotected	8.59 GB
A81 unprotected	8.59 GB	A01 protected	8.59 GB
A02 protected	17.55 GB	A82 unprotected	17.55 GB
A82 unprotected	17.55 GB	A02 protected	17.55 GB
A05 protected	35.17 GB	A85 unprotected	35.17 GB
A85 unprotected	35.17 GB	A05 protected	35.17 GB
A04 protected	70.56 GB	A84 unprotected	70.56 GB
A84 unprotected	70.56 GB	A04 protected	70.56 GB
A06 protected	141.12 GB	A86 unprotected	141.12 GB
A86 unprotected	141.12 GB	A06 protected	141.12 GB
A07 protected	282.35 GB	A87 unprotected	282.35 GB
A87 unprotected	282.35 GB	A07 protected	282.35 GB
050 unprotected	variable	099 protected	variable

Decimal gigabyte (GB) is 10^9 bytes.

Note: You must ensure that the volume data type is compatible with the host systems that can access this volume.

-type ess | ds | blocks
(Optional) Specifies the unit type of capacity that is specified by the **-cap** parameter.

ess Specifies that the unit is decimal gigabytes (GB) 10^9 bytes.

ds Specifies that the unit is gibibytes (GiB) 2^30 bytes.

blocks Specifies that the unit is 512 blocks.

TiB Specifies that the unit is tebibytes (TiB) 2^40 bytes. This value is an alias of ds.

GiB Specifies that the unit is gibibytes (GiB) 2^30 bytes.

MiB Specifies that the unit is mebibytes (MiB) 2^20 bytes.

Notes:

1. If the **-type** parameter is not specified, the lun is created as type **ds**.
2. The **-type** parameter is ignored when the **-os400** parameter is specified.

-cap *capacity*
(Optional) Specifies the storage size that is allocated to this volume object. The maximum volume size varies and depends on model and type.

Note: This parameter is required if the **-os400** parameter is not specified.

- If the **-type** parameter is omitted or the **-type ds** parameter is specified, the capacity value is the volume size in gibibytes (GiB), where one gibibyte (GiB) = 1 073 741 824 (2^30 bytes).
- If the **-type ess** parameter is specified, the capacity value is the volume size in gigabytes (GB), to the nearest 1/10 GB (format xxxx.x), where 1 GB = 1 000 000 000 (10^9 bytes).
- If the **-type blocks** parameter is specified, the capacity value is the volume size in 512-byte blocks.

-name *volume_name*
(Optional) Your nickname for this volume. The nickname can be 16 characters in length and can contain one of the following wildcard characters:

- #d decimal volume ID
- #h hexadecimal volume ID

-volgrp *volume_group_ID*
(Optional) Specifies to which volume group the volumes are assigned. A volume group ID is a four-digit decimal number with no leading zeros, prefixed with the letter V.

-wait
(Optional) Delays the command response until the volume configuration processes complete.

Note: If you specify this parameter, you must wait until your original command processes completely before you can enter a new command.

-sam standard | tse | ese
(Optional) Specifies the storage allocation method as follows:

standard
Designates that you want the system to fully allocate the volume with real extents when it creates the volumes. This value is the default.

tse Designates that you want the system to create track space-efficient volumes. After creation, these space-efficient volumes contain a set of virtual extents that are associated with the space-efficient storage in the same extent pool. The physical space for a given logical track on a track space-efficient logical volume is dynamically allocated and deallocated from the repository in the space-efficient storage.

To use this subparameter, you must have previously created space-efficient storage (by using the **mksestg** command) for the extent pool.

Note: Track space-efficient (TSE) repositories are required to create TSE volumes. However, starting with DS8000 Release 8, TSE repositories can no longer be created.

ese Designates that an extent space efficient (ESE) logical volume is provisioned with a set of virtual extents that are associated with the space efficient storage in the same extent pool. Physical space for an extent space efficient logical volume is dynamically allocated and de-allocated from the extent pool. ESE volumes are used for thin provisioning.

Note: To use this subparameter, you must have previously created space-efficient storage (by using the **mksestg** command) for the extent pool.

-eam rotateexts | rotatevols
(Optional) Specifies the extent allocation method as follows:

rotateexts
Specifies that the extents for each new logical volume are allocated across all available ranks, and is also known as storage-pool striping. This value is the default.

rotatevols
Specifies that the extents for each new logical volume are allocated from each successive rank. This means that the extents for a particular volume will be allocated from one rank, while the extents for the next volume will be allocated from the next successive rank, and so on.

-perfgrp *performance_group_ID*
(Optional) Specifies the performance group ID that the volumes are assigned to. The performance group ID begins with the letters *PG*. The default is *PG0*.

-resgrp *resource_group_ID*
(Optional) Specifies the resource group that the volumes are assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number. The default is *RG0*.

Note: If you create any fixed block LSSs with this command, the LSSs are assigned to the same resource group as the logical volumes.

-t10dif
(Optional) Specifies that the CRC-16-T10-DIF algorithm is being used to store data.

Note: The **-t10dif** and **-os400** parameters cannot be used together.

volume_ID ... | **-**
(Required) An array of one or more fixed block volume IDs to be created. The volumes must share a common logical subsystem ID.

The volume ID format is four hexadecimal characters *XYZZ* that represents the following values.

X Specifies the address group, 0–F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

To specify a range of volume IDs, separate the volume IDs with a dash (-).

You must separate multiple volume IDs or ranges of volume IDs with a blank space between each ID or range of IDs.

Example: 0100-010F 0180-018F 0120

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 1

Invoking the mkfbvol command

```
dscli> mkfbvol -dev IBM.2107-75FA120
-extpool P1 -name my_vol_#d -type ess -cap 8.6
-sam ese 0100 0101 0102 0103
```

The resulting output

```
FB volume 0100 successfully created.
FB volume 0101 successfully created.
FB volume 0102 successfully created.
FB volume 0103 successfully created.
```

Example 2

Invoking the mkfbvol command

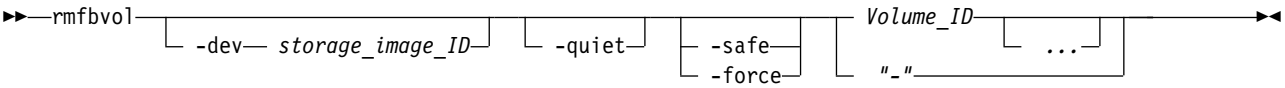
```
dscli> mkfbvol -extpool p0 -os400 050 -cap 5 0002
```

The resulting output

```
CMUC00025I mkfbvol: FB volume 0002 successfully created.
```

rmfbvol

The **rmfbvol** command deletes fixed block volumes from a storage image.



Parameters

Notes:

1. You can use this command when there are volumes that are in the configuration error state. To correct this configuration state, issue the **rmfbvol** command for each affected volume. You can specify a volume range according to the command specifications when it is appropriate.
2. Before Release 5.1, the DS8000 system did not check before deleting a volume. With Release 5.1 and later, the DS8000 system does not delete a volume that is in use. The phrase *in use* means that the volume is participating in a Copy Services relationship or is in a z/OS path group. Use the **-force** parameter to bypass the in-use checking and delete the volume. Use the **-safe** parameter for an additional check for volume participation in a non-default volume group before deleting the volume.

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID for all logical volumes, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-quiet
(Optional) Turns off the volume removal confirmation prompt for this command.

-safe
(Optional) The **-safe** parameter specifies that the system run a check to see whether the specified volumes are assigned to any non-default volume groups. If any volumes are still assigned to a user-defined volume group, the **rmfbvol** command fails without deleting any volumes. During this occurrence, messages are provided which list the volumes that are still assigned to a user-defined volume group.

Notes:

- If there is any reason that the system cannot run the check, the **rmfbvol** command fails and no volumes are deleted.
- The **-safe** and **-force** parameters cannot be specified at the same time.

-force
(Optional) The **-force** parameter allows for specified volumes to be deleted without checking to see

whether the volumes are in use. The phrase *in use* means that the volume is participating in a Copy Services relationship or is in a z/OS path group. If multiple volumes are specified and some are in use and some are not, the ones not in use can be deleted.

Note: The **-safe** and **-force** parameters cannot be specified at the same time.

Volume_ID ... | -
(Required) Specifies an array of one or more fixed block volume IDs to be removed. This parameter also accepts a fully qualified volume ID, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified.

Example of -dev parameter use

If you specify the **-dev** parameter, you can use the shortened version of the **Volume_ID** parameter as follows:

```
dscli> rmfbvol -dev IBM.2107-75FA120 0100-010F 0180-018F 0120
```

If you do not specify the **-dev** parameter and you specify the *Volume_ID* variable, you must use the fully qualified version of the volume ID as follows:

```
dscli> rmfbvol  IBM.2107-75FA120/0100-010F 0180-018F 0120
```

The volume ID is a 32 bit number that can be represented as four hexadecimal digits in the form of *XYZZ* where:

X Specifies the address group, 0–F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

To specify a range of volume IDs, separate the volume IDs with a hyphen.

You must separate multiple volume IDs or ranges of volume IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmfbvol command

```
dscli> rmfbvol -dev IBM.2107-75FA120 0100 0101
```

The resulting output

```
Are you sure you want to delete 0100? y/n Y
Volume 0100 successfully deleted.
Are you sure you want to delete 0101? y/n Y
Volume 0101 successfully deleted.
```

The following example shows the output that results when you specify that you want a range of volume IDs to be removed.

Invoking the rmfbvol command

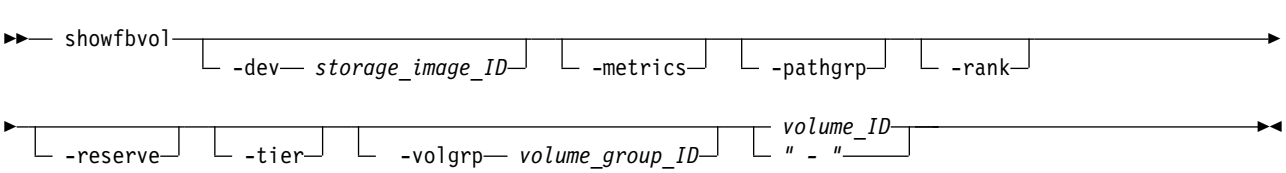
```
dscli> rmfbvol -dev IBM.2107-75FA120 0005-0024
```

```
CMUC00027W rmfbvol: Are you sure you want to delete FB
volume 0005-0024?
[y/n]:y
CMUC00028I rmfbvol: FB volume 0005 successfully deleted.
CMUC00028I rmfbvol: FB volume 0006 successfully deleted.
CMUC00028I rmfbvol: FB volume 0007 successfully deleted.
CMUC00028I rmfbvol: FB volume 0008 successfully deleted.
```

```
CMUC00028I rmfbvol: FB volume 0009 successfully deleted.
CMUC00028I rmfbvol: FB volume 000A successfully deleted.
CMUC00028I rmfbvol: FB volume 000B successfully deleted.
CMUC00028I rmfbvol: FB volume 000C successfully deleted.
CMUC00028I rmfbvol: FB volume 000D successfully deleted.
CMUC00028I rmfbvol: FB volume 000E successfully deleted.
CMUC00028I rmfbvol: FB volume 000F successfully deleted.
CMUC00028I rmfbvol: FB volume 0010 successfully deleted.
CMUC00028I rmfbvol: FB volume 0011 successfully deleted.
CMUC00028I rmfbvol: FB volume 0012 successfully deleted.
CMUC00028I rmfbvol: FB volume 0013 successfully deleted.
CMUC00028I rmfbvol: FB volume 0014 successfully deleted.
```

showfbvol

The **showfbvol** command displays detailed properties for an individual fixed block (FB) volume. Detailed information includes performance metrics of a fixed block (FB) volume and information about Safeguarded Copy capacity.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID. It is also required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-metrics
(Optional) Displays volume ID and performance metrics for the specified volume.

Notes:

1. All performance counts are an accumulation since the most recent counter wrap or counter reset. Volume performance counters are reset on a power-up sequence. Volume performance counters are reset by a server failover and failback sequence.
2. Do not use this parameter with the **-pathgrp**, **-rank**, **-reserve**, **-tier**, or **-volgrp** parameters.

-pathgrp
(Optional) Displays the path group status table, which contains path group information. This information includes the path group ID, the grouped, reserved, and path mode status.

Do not use this parameter with the **-metrics**, **-rank**, **-reserve**, **-tier**, or **-volgrp** parameters.

-rank
(Optional) Specifies that a rank extents table is to be displayed. This table displays the set of ranks that the logical volume has extents that are configured on and the number of extents for that logical volume.

Note: Do not use this parameter with the **-metrics**, **-pathgrp**, **-reserve**, **-tier**, or **-volgrp** parameters.

-reserve
(Optional) Displays the SCSI reserve status table, which contains SCSI reserve information. This information includes the WWPN, type, and the I/O port ID.

Do not use this parameter with the **-metrics**, **-pathgrp**,**-rank**, **-tier**, or **-volgrp** parameters.

-tier
(Optional) Displays the tier distribution table. The set of tiers are displayed with allocated storage for the specified logical volume and the percentage of the logical volume that is allocated on each tier.

Do not use this parameter with the **-metrics**, **-pathgrp**, **-rank**, **-reserve**, or **-volgrp** parameters.

-volgrp *volume_group_ID*
(Required if you do not specify the *volume_ID* parameter.) Specifies that the fixed block volumes that are associated with the designated volume group ID are to be displayed.

Notes:

1. You can use the **-volgrp** parameter only when you are doing a query for performance metrics.
2. Do not use the **-volgrp** parameter with the *volume_ID* parameter.
3. Do not use the **-volgrp** parameter with the **-metrics**, **-pathgrp**, **-rank**, **-reserve**, or **-tier** parameters.

volume_ID | –
(Required if you do not specify the **-volgrp** parameter.) Displays information for the specified volume. This parameter accepts a fully qualified volume ID, which consists of the *storage_image_ID* or a shortened version without the storage image ID, if you specify the **-dev** parameter. The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ* where:

X Specifies the address group, 0–F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) while you are in the DS CLI interactive command mode.

Note: Do not use the *volume_ID* parameter with the **-volgrp** parameter.

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output reports that are associated with the **showfbvol** command that uses the **-rank** parameter. When the rank parameter is specified, a rank extents table is also displayed. It appears at the end of the regular report.

Invoking the showfbvol command to show volume properties

Note: The example output is based on using the **showfbvol** command for a 1.0 gibibyte (GiB)volume. When the rank parameter is specified, a rank extents table is displayed at the end of the regular report.

```
dscli> showfbvol -dev IBM.2107-1300861 -rank 0000
```

The resulting output

Name	ID	acc state	data state	config state	device MTM	data type	addrgrp
RegrFB Vol1	0006	Online	Normal	Normal	2107-900	FB 512	0

extpool	exts	captype	cap (2^30B)	cap (10^9B)	cap (blocks)	volgrp	ranks	dbexts
P0	1	DS	1.0	-	2097152	-	1	0

sam	repcapalloc	eam	reqcap (blocks)
Standard	-	managed	2097152

realextents	virtualextents	migrating	perfgrp	migratingfrom	resgrp
1	2	0	PGO	-	RG0

etmonpauseremain	etmonitorreset	GUID
-	unknown	6005076303XXC0XD0000000000000000

safeguardedcap (2^30B)	safeguardedloc	usedsafeguardedcap (2^30B)	safeguarded
16.0	mypool_1	1.0	yes

```
============================Rank extents=================================
```

Rank	Extents	Capacity(MiB)
R0	1	1
R2	2	1

Report field definitions (*-metrics* parameter not specified)

Name
Indicates the nickname that you assigned for this volume object.

ID Indicates the unique identifier that is assigned to this volume object.

Accstate
One of the following access states are displayed: Online or Fenced.

Online
The logical volume is accessible to a host.

Fenced
The logical volume is in the volume fenced state and is not accessible to the host.

Datastate
One of the following data states are displayed:

Normal
None of the other data states apply. The access state is Online.

Pinned
Indicates that none of the other data states apply and the logical volume has one or more pinned non-retryable tracks. The access state is Online.

Read only
Indicates that the logical volume can be read but not written to because one or more extents on the logical volume are on a rank in the read only data state. The access state is Online.

Inaccessible
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the inaccessible data state. The access state is Fenced.

Virtual space fault
Indicates that the logical volume has a storage allocation method of extent space-efficient or track space-efficient. Not enough space was available to convert a virtual logical track to a real logical track. The access state is Online.

Indeterminate data loss
Indicates that the following data states do not apply and that one of the following conditions occurred:

Data states that do not apply:

- Rank failed
- Rank repairing
- Rank repaired
- Global inaccessible
- Global lost data

Conditions - one of the following conditions occurred:

- Committed write data was lost before it was de-staged and the track identifiers that are associated with the data are unknown.
- Data was lost that indicated extents on the logical volume were active FlashCopy targets.

The access state is Fenced.

Rank failed
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the Failed data state. The access state is Fenced. This data state changes to Rank repairing if the rank changes to the Rank repairing state through use of the repair array function.

Rank Repairing
Indicates that one or more extents that are associated with the logical volume are on ranks in the repairing data state. The access state is Fenced.

Rank Repaired
Indicates that one or more extents that are associated with the logical volume are on ranks that were in the repairing state, but are not in the repairing state now. The access state is Fenced.

Global inaccessible
Indicates that the global metadata that is associated with the logical volume configuration is inaccessible. Some of the data that is associated with the logical volume might be inaccurate. The access state is Fenced.

Global lost
Indicates that global metadata that is associated with the logical volume configuration was lost. As a result, some of the data that is associated with the logical volume might be inaccurate. The access state is Fenced.

NVS data inaccessible
Indicates that active nonvolatile storage (NVS) data is inaccessible for one or more logical volumes of an LSS group. The logical volumes in the LSS group cannot be made accessible. The access state is Fenced.

Configstate
One of the following configuration states are displayed:

Normal
Indicates that no logical volume configuration operations are in progress, and the volume is not being deconfigured, merged, or migrated.

Configuring
Indicates that the logical volume is being configured for the first time.

Reconfiguring
Indicates that the logical volume is allocating or de-allocating extents due to a modification of the requested capacity attribute after initial creation.

Deconfiguring
Indicates that the logical volume is being deleted.

Configuration error
Indicates that the initial configuration did not complete successfully. This state reflects an internal error condition and not an error in the request to create the volume. If you have a volume in this state, use the **rmfbvol** command to delete each volume that is listed with the configuration state "configuration error".

Merging
Indicates that the volume is being merged. For example, merging from one extent pool to a different extent pool.

Migrating
Indicates that the volume is migrating, or waiting to be migrated.

Migration Cancelled
Indicates that the volume was migrating and then the 'migcancel' action of the **manageckdvol** command was issued. This process left some extents waiting to be migrated in the source pool while other extents were migrated to the target pool. Migration stopped, and cannot be resumed. If you have a volume in this state, try to migrate it again to the original source or target extent pool.

Migration Paused
Indicates that the volume was migrating and then the 'migpause' action of the **manageckdvol** command was issued. Migration stopped, but can be resumed.

Migration Error
Indicates that the volume migration process failed to complete successfully. This state reflects an internal error condition and not an error in the user's request to migrate a volume. If you have a volume in this state, try to migrate it again to the original source or target extent pool.

Reconfiguration error
Indicates that the reconfiguration request did not complete successfully.

Deconfiguration error
Indicates that a request to delete a volume did not complete successfully. This state reflects an internal error condition and not an error in the request to remove the volume. To correct this state, you must reissue the **rmfbvol** command for the designated volume.

Transposition Error
Indicates that the volume is in an extent pool that was unsuccessfully merged. This state reflects an internal error condition. **Corrective action**: Use the **chextpool** command with the **-merge** parameter again to redrive the merge extent pool and to correct this state.

device MTM
Indicates the volume device type and the machine type. The volume MTM value is determined by the fixed block volume data type and the provisioned capacity (in GB). The machine type is either 2107 or 1750; however, the MTM value can be any one of the following depending on your system:

2107-900
Indicates a standard 2107 volume.

1750-500
Indicates a standard 1750 volume.

xxxx-A0*x*
The *xxxx* is 2107 or 1750; the A0 indicates a System i protected volume (for example, 2107-A01 or 1750-A07).

xxxx-A8*x*
The *xxxx* is 2107 or 1750; the A8 indicates a System i unprotected volume (for example, 2107-A81 or 1750-A87).

Datatype
Indicates the volume data type setting. One of the following values is displayed:

- FB 512
- FB 512T
- FB 520P
- FB 520U

Addrgrp
Indicates the address group that contains the designated volume object. An address group ID is one hexadecimal character (0 - F).

Extpool
Indicates the extent pool ID. Volume extents are allocated from this extent pool ID.

Note: Volumes that belong to an encrypted extent pool are encrypted. You can see the key group of an extent pool by using the **lsextpool -l**, or **showextpool** commands.

Exts
Indicates the number of real and virtual extents that are used by the designated volume ID.

Captype
Indicates capacity unit type that is used at volume creation. One of the following values is displayed:

ESS
The capacity unit is decimal gigabytes (GB).

DS The capacity unit is gigibytes (GiB).

DS/ESS
The capacity unit is gigibytes (GiB) or decimal gigabytes (GB).

Blocks
The capacity unit 512 B.

iSeries
The capacity unit was not specified at volume creation. This fixed block volume was created for IBM Z systems.

cap (MiB)

The size of volume that is available for host system access in MiB(2^{20}B).

Cap (2^30B)
Indicates the size of the volume that is available for host system access in gibibytes GiB(2^{30}B).

Note: " – " is displayed if the capacity unit type of the volume is ESS (captype=ESS)

Cap (10^9B)
Indicates the size of volume that is available for host system access in decimal gibibytes GiB(10^{9}B).

Note: " – " is displayed if the capacity unit type of the volume is DS (captype=DS)

Cap blocks
Indicates the quantity of volume logical blocks that are available for host system access.

Volgrp
Indicates the volume groups (excluding default volume groups) that a volume belongs to.

Multiple volume groups that are associated with the volume are separated by a comma.

A " – " displays if no volume groups exist that are associated with the volume.

`Unknown` displays if information about the volume groups is not available.

Ranks
Indicates the number of ranks that the volume resides on.

SAM
Indicates the storage allocation method. The following values are displayed:

standard
Designates that the system fully allocated the volume with real extents at volume creation time.

tse
Designates that a track space-efficient logical volume contains a set of virtual extents that are associated with the space-efficient storage in the same extent pool. Physical space for a given logical track on a track space-efficient logical volume is dynamically allocated and deallocated from the repository in the space-efficient storage.

Note: Track space-efficient (TSE) repositories are required to create TSE volumes. However, starting with DS8000 Release 8, TSE repositories can no longer be created.

ese
Designates that an extent space efficient logical volume is provisioned with a set of virtual extents that are associated with the space efficient storage in the same extent pool. Physical space for an extent space efficient logical volume is dynamically allocated and deallocated from the extent pool.

Repcapalloc
Indicates the allocated physical repository capacity of the track space-efficient storage. This value is calculated on the available repository capacity as a result of writes to the track space-efficient volume. This value is displayed in the format of X.Y, where X is in whole gibibytes (GiB) and Y represents tenths of a GiB, which is limited to a single digit (0 - 9). .

Note: A " – " value is displayed in this column.

EAM
Indicates the extent allocation method that is to be used if the volume is migrated or expanded. One of the following values is displayed:

legacy
Designates that the volume was created before the use of the current algorithm.

rotateexts
Indicates that the extents for each new logical volume are allocated across all available ranks, and is also known as storage-pool striping. This value is the default.

rotatevols
Indicates that the extents for each new logical volume are allocated from each successive rank. This means that the extents for a particular volume will be allocated from one rank, while the extents for the next volume will be allocated from the next successive rank, and so on.

managed
Indicates that the extents are currently managed by Easy Tier, and the extents for any new volumes are initially allocated across all available ranks in the lowest tier of storage.

" – "
The value " - " is displayed if the extent allocation method does not apply, for example, track space-efficient logical volumes.

Reqcap (blocks)
Indicates the requested quantity of volume logical block (for example, 3339).

realextents
Indicates the number of real extents that are used by the logical volume.

virtualextents
Indicates the number of virtual extents that are used by the logical volume.

realcap(MiB)
Indicates the number of real capacity that is used by the logical volume in MiB(2^{20}B.

migrating
The number of extents for this volume that are currently being migrated.

migratingcap(MiB)
The capacity for this volume that is being migrated in MiB(2^{20}B.

migratingfrom
A list of one or more extent pool IDs where the extents are migrating from. If no migrating extents exist, a dash "-" is displayed. `Unknown` is displayed if information about the extent pool IDs is not available.

perfgrp
Indicates the performance group ID that the volume is assigned to. The performance group ID begins with the letters *PG* and ends with a decimal number.

resgrp
Indicates the resource group ID that the volume is assigned to. The resource group ID begins with the letters *RG* and ends with a decimal number.

tierassignstatus
Status of assigning a volume to a target tier.

Assign Pending
An assign action was specified, but has not started.

Assign Pending Hardware
An assign action was specified, but not started because of a hardware condition.

Assigning
An assign action is in progress.

Assigned
The volume was assigned to the specified tier.

Unassign Pending
An unassign action was specified, but has not started.

Error An assign action failed. See the **tierassignerror** value for the reason.

Unknown
An assign action was specified, but the specific action is unknown.

" – " No assign action was specified (none).

tierassignerror
Failure reason if assign action status is Error.

Easy Tier not active
Easy Tier is not active. See the `etmanaged` column from **lsextpool** to view whether the volume is in a pool that is managed by Easy Tier.

Use `manageckdvol -action tierunassign` to unassign the volume, ensure that the pool is managed by Easy Tier (see **chsi**), and then use `manageckdvol -action tierassign` to assign the volume again.

Target Tier not available
The specified target tier does not currently exist. Use `manageckdvol -action tierunassign` to unassign the volume, ensure that space is available on the specified tier, and then use `manageckdvol -action tierassign` to assign the volume again.

Tier definitions have changed
The target tier was valid, but defined tiers changed internally and the target tier is no longer valid. Use `manageckdvol -action tierunassign` to unassign the volume, and then use `manageckdvol -action tierassign` to assign the volume again.

" – " The assign status is not Error.

tierassignorder
The method that is used to define the assigning order.

Access
Assign extents only when accessed.

ETdata
Assign high usage extents first, based on Easy Tier data.

Unknown
Unknown assigning order method.

" – " No assign action was specified.

tierassigntarget
Assign action target tier.

FlashTier0
The tier of high performance flash drives.

FlashTier1
The first tier of high capacity flash drives.

FlashTier2
The second tier of high capacity flash drives.

ENT Enterprise tier that consists of drives with speeds of 10K RPM, 15K RPM, or a mixtures of 10K RPM and 15K RPM speeds.

NL Nearline tier consists of high-volume disks that are either SATA or SAS Nearline drives.

HDDExclude
Any tier except Enterprise (ENT) or Nearline (NL).

NLExclude
Enterprise tiers but not a Nearline tier.

" – " No assign action was specified.

%tierassigned
The percentage of the provisioned capacity that is assigned. The value is 0 (zero) if the volume is not assigned to a tier.

etmonpauseremain
Specifies the pause in Easy Tier monitoring. One of the following values is displayed:

0H1M-168H0M
Specifies the time (in hours and minutes) that remains of the pause in the Easy Tier monitoring process.

infinite
Specifies that Easy Tier monitoring remains paused until a resume action is submitted.

- The dash (-) specifies that Easy Tier monitoring is not paused.

unknown
Specifies that the system failed to query the time that remains of the pause.

etmonitorreset
Easy Tier extent pool monitoring reset date is as follows.

date
Specifies the date of the last Easy Tier monitoring reset in ISO 8601 format: `yyyy-MM-dd'T'HH:mm:ssZ`, where

yyyy Year

MM Month (01-12)

dd Day (01-31)

T Single letter T without quotes

HH Hour (00-23)

mm Minutes (00-59)

ss Seconds (00-59)

Z Time zone offset from UTC [-HHmm | +HHmm]

unknown
Specifies that the date in which Easy Tier monitoring of this extent pool was last reset is not known.

unsupported
Specifies that Easy Tier extent pool management is not supported.

GUID
Specifies the GUID (Globally Unique Identifier): The NAA (Network Addressing Authority) unique value that identifies a volume object.

safeguardedcap(2^30B)
Specifies the amount of safeguarded backup capacity that is reserved for safeguarded copies.

safeguardedloc
Specifies the pool where the safeguarded copies are written to.

usedsafeguardedcap(2^30B)
Specifies the amount of physical pool capacity that is used by the safeguarded copies.

safeguarded
Specifies whether Safeguarded Copy is enabled for the specified volume.

Report field definitions (*-rank* parameter specified)

Rank (Rank Extent table)
Indicates the rank ID.

Extents (Rank Extents table)
Indicates the number of extents for the volume on the rank.

Capacity(MiB)(Rank Extents table)
Indicates the capacity for the volume on the rank in MiB (2^{20}B).

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output reports that are associated with the **showfbvol** command using the **-metrics** parameter.

Invoking the showfbvol to show performance metrics

```
dscli> showfbvol -metrics IBM.2107-75FA120/0101
```

The resulting output

ID	Date	norm rdrqts	norm rdhits	norm write req	norm write hits	seq read reqs	seq read hits	seq write req
IBM.2107-75FA120/0101	10/11 /04 02:23:49	10000	10000	10000	10000	10000	10000	10000

seqwrite-hits	cachfwr-reqs	cachfwr-hits	cachfw-reqs	cachfw-hits	inbcach-load	bypass-cach	seq DASD trans
10000	10000	10000	10000	10000	10000	10000	10000

DASD-trans	cache-trans	NVS-spadel	norm write ops	seqwrite-ops	rec cache mis	qwrite-prots	CKDir-trkac
10000	10000	10000	10000	10000	10000	10000	0

CKD irtrk hits	cachsp-delay	timelow-ifact	phread	phwrite	phwrite	phbyte-read	phbyte-writ
0	10000	10000	10000	10000	10000	10000	10000

recmo-reads	sfile trk reads	contam-wrts	PPRC-trks	NVS-spallo	time-phread	timeph-write	byte-read
10000	0	0	10000	10000	10000	10000	10000

bytewrit	timeread	timewrite	zHPFRead	zHPFWrite
10000	10000	10000	-	-

zHPFPrefetchReq	zHPFPrefetchHit	GMCollisions-SidefileCount	GMCollisions-SendSyncCount
0	0	0	0

Report field definitions (*-metrics* parameter specified)

ID Indicates the unique identifier that is assigned to this volume object.

Date
Indicates the current time stamp for the volume performance counters.

normrdrqts
Indicates Search/Read Normal I/O Requests.

normrdhits
Indicates Search/Read Normal I/O Requests instances.

normwritereq
Indicates Write Normal I/O Requests.

normwritehits
Indicates DASD Fast Write I/O Request instances.

seqreadreqs
Indicates Search/Read Sequential I/O Requests.

seqreadhits
Indicates Search/Read Sequential I/O Request instances.

seqwritereq
Indicates Write Sequential I/O Requests.

seqwritehits
Indicates DASD Fast Write Sequential I/O Request instances.

cachfwrreqs
Indicates Search/Read Cache Fast Write I/O Requests.

cachfwrhits
Indicates Search/Read Cache Fast Write I/O Request instances.

cachfwreqs
Indicates Cache Fast Write I/O Requests.

cachfwhits
Indicates Cache Fast Write I/O Requests instances.

inbcachload
Indicates Inhibit Cache Loading I/O Requests that operate with DASD.

bypasscach
Indicates Bypass Cache I/O Requests.

seqDASDtrans
Indicates Sequential DASD to Cache Transfer Operations.

DASDtrans
Indicates DASD to Cache Transfer Operation Count.

cachetrans
Indicates Cache to DASD Transfer Operation Count.

NVSspadel
Indicates DASD Fast Write Operations Delayed Due to nonvolatile storage Space Constraints.

normwriteops
Indicates Normal 'DASD Fast Write' Write Operation Counts.

seqwriteops
Indicates Sequential Access 'DASD Fast Write' Write Operation Counts.

reccachemis
Indicates Number of record cache Read Misses.

qwriteprots
Indicates Quick Write Promotes.

CKDirtrkac
Indicates Irregular Track Accesses. The value 0 (zero) is displayed for a fixed block volume.

CKDirtrkhits
Indicates Irregular Track Accesses instances. The value 0 (zero) is displayed for a fixed block volume.

cachspdelay
Indicates Operations Delayed Due To Cache Space Constraints.

timelowifact
Indicates Milliseconds of lower interface I/O activity for the indicated device.

phread
Indicates Physical Storage Read Operations.

phwrite
Indicates Physical Storage Write Operations.

phbyteread
Indicates Physical Storage Bytes Read in 128 KB increments.

phbytewrit
Indicates Physical Storage Bytes Written in 128 KB increments.

recmoreads
Indicates Record Mode Read Operations.

sfiletrkreads
Indicates the number of tracks that are read from the Concurrent Copy or XRC Sidefile. The value 0 (zero) is displayed for a fixed block volume.

contamwrts
Indicates the Number of Contaminating writes for a Concurrent Copy or XRC volume. The value 0 (zero) is displayed for a fixed block volume.

PPRCtrks
Indicates the Number of tracks or portion of tracks that were transferred to the secondary device of a PPRC pair.

NVSspallo
Indicates the NVS Space Allocations.

timephread
Indicates the Physical Storage Read Response Time in 16 ms increments.

timephwrite
Indicates the Physical Storage Write Response Time in 16 ms increments.

byteread
Indicates the number of Bytes read in 128 KB increments.

bytewrit
Indicates the number of Bytes written in 128 KB increments.

timeread
Indicates the accumulated response time for all read operations.

timewrite
Indicates the accumulated response time for all write operations.

zHPFRead
Indicates the HPF Read I/O Requests for volume performance statistics.

zHPFWrite
Indicates the HPF Write I/O Requests for volume performance statistics.

zHPFPrefetchReq
Indicates the number of HPF Pre-fetch I/O requests.

zHPFPrefetchHit
Indicates the number of HPF Pre-fetch I/O request hits.

GMCollisionsSidefileCount
Indicates the number of Global Mirror Collisions sidefile.

GMCollisionsSendSyncCount
Indicates the number of Global Mirror Collisions Send Synchronous Count.

Example

Show tier statistics

```
dscli> showfbvol –tier 0000
```

The resulting output

```
Name              myvol0800
ID                0800
accstate          Online
datastate         Normal
configstate       Normal
...
migrating         20
perfgrp           PG0
migratingfrom     P0
resgrp            RG1
tierassignstatus  Assigning
tierassignerror   -
tierassignorder   ETdata
tierassigntarget  FlashTier0
%tierassigned     54
etmonpauseremain  1H44M
etmonitorreset    2013-07-26T14:00:00+07
============== Tier Distribution =================
Tier %allocated
=====================
FlashTier0     54
ENT            46
```

Report field definitions (*-tier* parameter is specified)

Tier Tier ID

FlashTier0
The tier of high performance flash drives.

FlashTier1
The first tier of high capacity flash drives.

FlashTier2
The second tier of high capacity flash drives.

ENT Enterprise tier, which consists of drives with speeds of 10K RPM, 15K RPM, or a combination of drives of both speeds.

NL Nearline tier, which consists of high volume SATA or SAS Nearline disk drives.

Unknown
Tier is unknown.

%allocated
The percentage of provisioned capacity on this tier.

Example 4

Specifying the -pathgrp parameter

If you specify the -pathgrp parameter and no path groups exist for this volume, the following message is displayed.

```
CMUC00234I lsfbvol: No Path Groups found.
```

If you specify the **-pathgrp** parameter and path groups exist for this volume, a path group status table is appended to the resulting output.

```
dscli> showfbvol –pathgrp efff
```

The resulting output

```
Name           efff
ID             EFFF
accstate       Online
datastate      Normal
configstate    Normal
...
migrating      0
perfgrp        PG31
migratingfrom  -
resgrp         RG62
=========================Path Group status=======================
GroupID                  State       Reserve    Mode        Sysplex
=================================================================
800002B9472827CA78BC17   Ungrouped   Disabled   Single      N/A
880005B9472827CAAD6FBA   Grouped     Disabled   Multi-path  N/A
800009B9472827CAC684B9   Ungrouped   Disabled   Single      PLEXM1
```

Report field definitions (*-pathgrp* parameter is specified)

GroupID
The path group ID. An 11-byte value that is displayed as 22 hexadecimal characters.

Note: The path group ID is supplied by the host and is not interpreted further by the DS8000 system. This means that the hosts are free to define, or redefine, the meaning of this value with no impact to the DS8000 system. However, some programs such as ICKDSF break drown the ID into distinct fields with the following partial display of the ICKDSF logical path status table.

```
+--------------------------+-
|       PATH GROUP ID      | ...
|------+------+----+--------+-
|      |      |CPU |CPU TIME|
|  ID  |SERIAL|TYPE|  STAMP | ...
+------+------+----+--------+-
|800002|B947  |2827|CA78BC17| ...
+------+------+----+--------+-
|880005|B947  |2827|CAAD6FBA| ...
+------+------+----+--------+-
|800009|B947  |2827|CAC684B9| ...
+------+------+----+--------+-
```

For more information, see ICKDSF User's Guide and Reference.

State The grouped state of this path group. Valid state values are “Grouped” or “Ungrouped”.

Reserve
The reserved state of this path group. Valid state values are “Enabled” or “Disabled”.

Mode The path mode for this path group. Valid mode values are “Single” or “Multi-path”.

Sysplex
The z/OS sysplex name. If the name is not set or available, N/A is displayed.

Example 5

Specifying the -reserve parameter

If you specify the **-reserve** parameter and there are no SCSI reservations for this volume, the following message is displayed.

```
CMUC00234I lsfbvol: No SCSI reservations found.
```

If you specify the **-reserve** parameter and SCSI reservations exist for this volume, the SCSI reservation attributes and a SCSI reserve port table is appended to the resulting output.

```
dscli> showfbvol -reserve 0200
```

The resulting output

```
ID              0200
accstate        Online
datastate       Normal
configstate     Normal
...
migrating       0
perfgrp         PG0
migratingfrom   -
resgrp          RG0
========SCSI Reserve Status========
PortID WWPN               ReserveType
==================================
I0040  500507630310003D   Persistent
I0041  500507630310403D   Persistent
-      50050763080805BB   Persistent
-      50050763080845BB   Persistent
```

Report field definitions (*-reserve* parameter is specified)

PortID
The I/O port ID. If the host is online, then the I/O port ID is displayed and is formatted as a leading uppercase letter "I" followed by four hexadecimal characters (for example, I0040). If the host is not online, the field contains a ‘-‘ (dash).

WWPN
The worldwide port name that is displayed as 16 hexadecimal characters.

ReserveType
The SCSI reservation type for all connections. Valid reservation types are “Traditional”, “Persistent”, or “PPRC”.

Volume group specific commands

Volume groups require the use of specific commands that are used to create, modify, and delete volume groups and to display volume group information.

The following volume group specific commands are available:

chvolgrp
Modifies a volume group name and volume members.

lsvolgrp
Generates a report that displays a list of volume groups in a storage system and the status information for each volume group in the list.

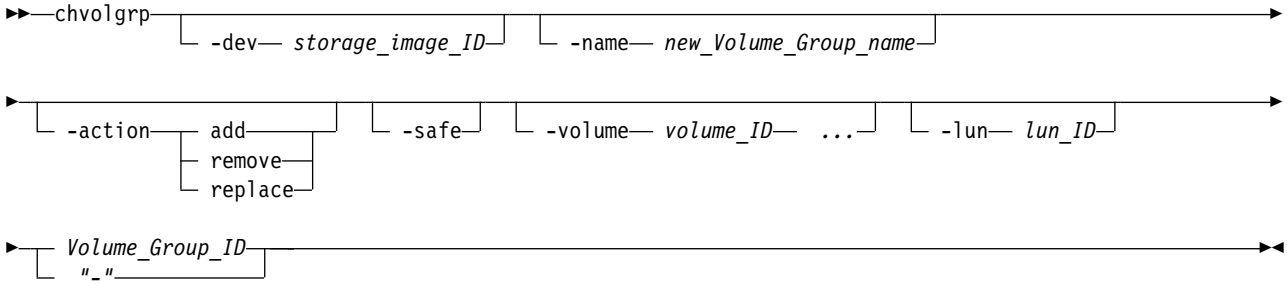
mkvolgrp
Creates a volume group in a storage system.

rmvolgrp
Deletes the specified volume group or volume groups from a storage system.

showvolgrp
Generates a report that displays the detailed properties of a volume group.

chvolgrp

The **chvolgrp** command modifies a volume group name and volume members.



Parameters

Note: If you are using an HP-UX operating system, the number of volumes in the volume group must not exceed seven volumes. This restriction applies only when the *hostconnect* attribute for the **-addrdiscovery** parameter is set to *reportlun* and the associated volume group is of type *mapscsi256*.

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify fully qualified IDs or do not set the *devid* variable in your profile or through the **setenv** command. The storage image ID is also required if the HMC is aware of more than one storage image. The **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-name *new_Volume_Group_name*
(Optional). Specifies a new name for the volume group. The name is limited to 16 characters. The name must be unique across volume groups that are contained by a storage unit.

-safe
(Optional). Enables checks for the safe removal of the volume from the volume group.

-action add | remove | replace
(Optional, unless the **-volume** parameter is specified). Specify one of the following values with this parameter:

add Specifies that the volumes be added to the volume group.

remove
Specifies that the volumes be removed from the volume group.

replace
Specifies that the existing volumes be replaced by the specified volumes.

Note: The **chvolgrp** command fails if you specify the **-volume** parameter and not included the **-action** parameter.

-safe
(Optional). Valid for the remove and replace actions, only. Enables checks for the safe removal or replacement of the volume from the volume group. The checked conditions that give a failed result include: an open host reservation, a Copy Services relationship, or an I/O received within the last 5 minutes.

-volume *volume_ID ...*
(Optional unless you are specifying the **-action** or the **-lun** parameter, then the **-volume** parameter is required.) Specifies an array of one or more volume IDs or volume ID ranges to be included in the volume group when the **-action** parameter is specified.

The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of *XYZZ* where:

X Specifies the address group, 0–F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

To specify a range of volume IDs, separate the volume IDs with a dash (-).

You must separate multiple volume IDs or ranges of volume IDs with a comma between each ID or range of IDs.

Notes:

1. For SCSI MAP 256, the array or ranges cannot exceed 256 volume ID entries. Otherwise, up to 64 384 entries are allowed.
2. The **chvolgrp** command fails if you specify the **-volume** parameter and do not specify the **-action** parameter.

Example: 0100-010F,0180-018F,0120

-lun *lun_ID*
(Optional - SCSI MAP 256 only). Specifies the LUN ID in hexadecimal value (00 - FF), which is mapped to the specified volume ID when the **-action** *add* or **-action** *replace* parameter is specified. If multiple volume IDs are specified by the **-volume** parameter, the LUN ID is consecutively assigned in incremental order. If the specified LUN ID is not valid, the command is rejected.

Note: This parameter is only valid when the target volume group type is SCSI MAP 256. Otherwise, this command fails.

If the **-action** *add* parameter is specified and the specified LUN ID is already mapped to the other volume in the specified volume group, the command fails.

If the **-action** *add* parameter is specified without the **-lun** parameter, an unused LUN ID is assigned to the volume ID. In this case, the unused LUN ID is selected from a smaller number.
The following example shows how this process works:

```
A volume group of "SCSI Map 256" type has Volume 0000 and 0001.
Their LUNs are the members of the following volume group:
(showvolgrp displays the current mapping.)
0000 : 10
0001 : 11
Because the range of LUN IDs is 00-FF, the unused LUN IDs
are 00,01,...,0F,12,13,...,FF.

If you add volume 0002 and 0003 to this volume group without
the -lun parameter, the mapping results in the following
because 00 and 01 are "smaller" unused LUN IDs:
```

```
0002 : 00
0003 : 01
0000 : 10
0001 : 11
```

If the **-action** *replace* parameter is specified without specifying the **-lun** parameter, lun_ID=00 is assumed.

Volume_Group_ID | –
(Required). Specifies the ID of the volume group that being changed. The volume group ID is made up of the storage image ID followed by the volume group ID. This parameter also accepts a fully qualified volume group ID including the storage image ID or a shortened version. The shortened version is a four-digit decimal number with no leading zeros, prefixed with the letter *V*.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the chvolgrp command

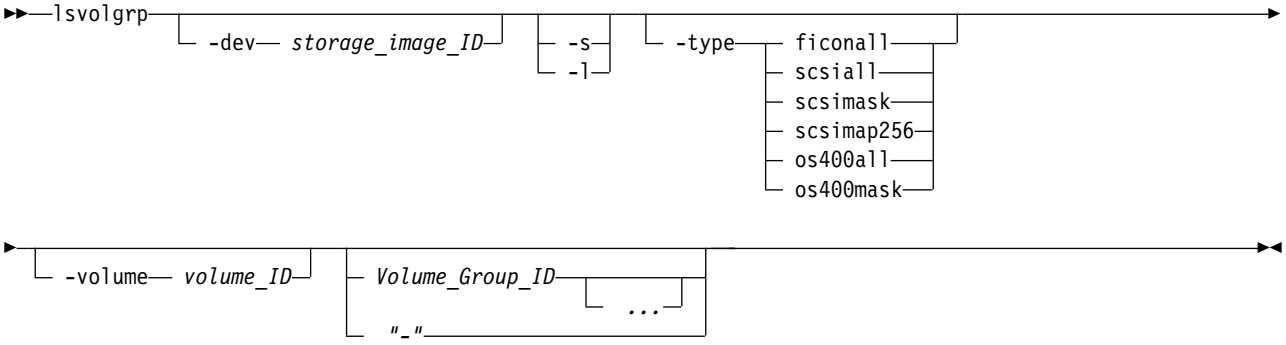
```
dscli> chvolgrp -action add
-volume 0000-000F IBM.2107-75FA120/V2341
```

The resulting output

```
Volume group V2341 successfully modified.
```

lsvolgrp

The **lsvolgrp** command displays a list of volume groups in a storage image and status information for each volume group in the list.



Parameters

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which includes manufacturer, machine type, and serial number. Displays only the objects for the storage unit that is specified. The storage image ID is required if you do not specify fully qualified IDs, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-s
(Optional). Displays volume group IDs only. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional). Displays the default output. You cannot use the **-l** and the **-s** parameters together.

-type | ficonall | scsiall | scsimask | scsimap256 | os400all | os400mask
(Optional). Displays only volume groups that are configured as the specified volume group type.

-volume *volume_ID*
(Optional). Displays volume groups that contain the specified volume ID. The volume ID is a 32 bit number that can be represented as four hexadecimal digits in the form of *XYZZ* where:

X Specifies the address group, 0-F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

Volume_Group_ID ... | -
(Optional). Displays volume groups with the specified IDs. A volume group ID is a four-digit decimal number with no leading zeroes, prefixed with the letter *V*.

This parameter accepts a fully qualified volume group ID or a shortened version, if the **-dev** parameter is specified.

To specify a range of volume group IDs, separate the volume group IDs with a hyphen. You must separate multiple volume group IDs or ranges of volume group IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lsvolgrp** command using the **-l** parameter.

Invoking the lsvolgrp command

```
dscli> lsvolgrp -dev IBM.2107-75FA120 -l
```

The resulting output

```
Tues Aug 11 02:23:49 PST 2016 IBM DS CLI Device: IBM.2107-75FA120
```

Name	ID	Type
Host_xxx _volumes	1011	OS400 Mask
Host_yyy _volumes	1111	OS400 Map 256
Host_zzz _volumes	1211	SCSI Mask

Report field definitions

Name
Indicates the name you assigned for this volume group ID.

ID*
Indicates the storage unit ID followed by the volume group ID. The volume group identifier is a four-digit decimal number, with no leading zeros, prefixed by the letter V.

Type
Indicates the type of the volume group.

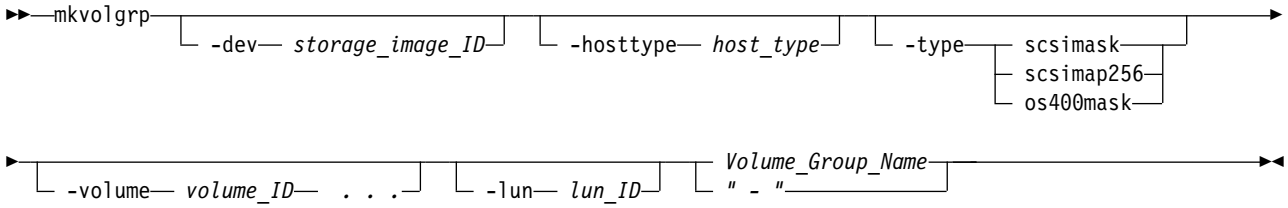
Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

mkvolgrp

The **mkvolgrp** command creates a volume group in a storage image.



Parameters

Notes:

1. The ID is assigned during volume group creation based on the current configuration, past configuration changes, and other internal considerations. No algorithm is available to accurately predict the newly created volume group ID in every case.
2. Ensure that you use the **-hosttype** parameter when you issue this command.

If you are using an HP-UX operating system, the number of volumes in the volume group must not exceed 7 volumes. This restriction only applies when the **hostconnect** attribute for the **-addrdiscovery** parameter is set to **reportlun** and the associated volume group is of type *scsimap256*.

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which includes manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the *-dev* parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-hosttype *host_type*
(Optional) Use this parameter as an alternative method for specifying the type of Volume Group.

Note: You cannot use this parameter with the **-type** parameter.

-type | scsimask | scsimap256 | os400mask
(Optional). Specifies the type of the volume group.

scsimask (default)
Creates a SCSI mask volume group. This option is available if the host adapter supports four-byte LUN addresses.

scsimap256
Creates a SCSI-MAP 256 volume group.

os400mask
Creates an OS/400 mask volume group. The IBM i host system typically uses fixed block volumes of 520-byte logical block size. This option is available only if the host adapter supports four-byte LUN addresses.

Note: This volume group is also referred to as SCSI520-MASK. When an error message is displayed for the os400mask, SCSI520-MASK is referenced instead.

Note: You cannot use this parameter with the **-type** parameter.

-volume *volume_ID* | **. . .**
(Optional). Specifies the array of volume IDs to include in the volume group. For the **-type** *scsimap256* parameter, the array cannot exceed 256 volume ID entries. Otherwise, up to 64 384 entries are allowed.

The volume ID is a 32-bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ* where:

X Specifies the address group, 0–F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

To specify a range of volume IDs, separate the volume IDs with a dash (–).

You must separate multiple volume IDs or ranges of volume IDs with a comma between each ID or range of IDs.

Example: 0100-010F,0180-018F,0120

-lun *lun_ID*
(Optional) Specifies the LUN ID in hexadecimal value (00 - FF) which is mapped to the specified volume ID for a SCSI-MAP256 type volume group. If multiple volume IDs are specified by the **-volume** parameter, LUN IDs are assigned consecutively in incremental order.

Note: This parameter is only valid for a SCSI-MAP 256 type volume group. If this parameter is specified for any other type of volume group, the command fails.

Volume_Group_Name | **-**
(Required). Specifies the volume group name, not to exceed 16 characters. Ensure that the name is unique within the scope of the storage image. Accepts a fully qualified volume group name or a shortened version, if the **-dev** parameter is specified.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the mkvolgrp command

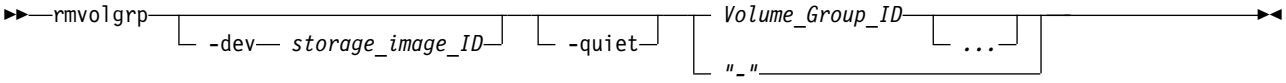
```
dscli> mkvolgrp -dev IBM.2107-75FA120 -volume 0000-000F host_xyz_volumes
```

The resulting output

```
Volume group V0 successfully created.
```

rmvolgrp

The **rmvolgrp** command deletes existing volume groups from a storage image.



Parameters

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-quiet
(Optional) Turns off the volume group removal confirmation prompt for this command.

Volume_Group_ID ... | -
(Required). Specifies an array of one or more volume groups IDs to be deleted. All volume groups specified must belong to the same storage unit. This parameter also accepts a fully qualified volume group ID, which consists of the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is a four-digit decimal number with no leading zeroes, prefixed with the letter *V*.

To specify a range of volume group IDs, separate the volume group IDs with a dash (-).

You must separate multiple volume group IDs or ranges of volume group IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example of -dev parameter use
If you specify the **-dev** parameter, you can use the shortened version of the *Volume_Group_ID* parameter as follows:

```
dscli> rmvolgrp -dev IBM.2107-75FA120 V11
```

If you do not specify the **-dev** parameter and you specify the *Volume_Group_ID* parameter, you must use the fully qualified version of the volume group ID as follows:

```
dscli> rmvolgrp IBM.2107-75FA120/V11
```

Example

Invoking the rmvolgrp command

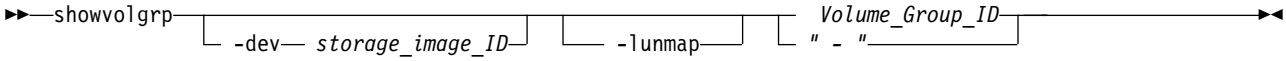
```
dscli> rmvolgrp IBM.2107-75FA1243/V123
```

The resulting output

```
Are you sure you want to delete Volume Group IBM.2107-75FA1243/V123? y/n
Y
Volume group IBM.2107-75FA1243/V123 successfully deleted.
```

showvolgrp

The **showvolgrp** command displays detailed properties of a volume group.



Parameters

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume group ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

Example: IBM.2107-75FA120

-lunmap
(Optional). Specifies that a LUN mapping table be displayed that shows the volume ID and LUN ID for all volumes in the specified volume group. This parameter is valid for all scsi and os400 type volume groups.

-etccmap
(Optional). Specifies that a LUN mapping table be displayed that shows the volume ID, LUN ID, and the Easy Tier cooperative caching status for all volumes in the specified volume group. While always valid, this parameter is meaningful only when the volumes in the volume group are managed by a host that has the Easy Tier Server utility installed.

Note: Starting with DS8900F, Release 8, Easy Tier Server is no longer supported.

Volume_Group_ID | **-**
(Required). Specifies that the properties be displayed for the specified volume group. This parameter accepts a fully qualified volume group ID, which consists of the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is a four-digit decimal number with no leading zeros, prefixed with the letter *V*.

Examples of -dev parameter use

If you specify the **-dev** parameter, you can use the shortened version of the *Volume_Group_ID* parameter as follows:

```
dscli> showvolgrp -dev IBM.2107-75FA120  V11
```

where *V11* represents value for the volume group ID.

If you do not specify the **-dev** parameter, and you specify the *Volume_Group_ID* parameter, you must specify the fully qualified version of the *Volume_Group_ID* parameter as follows:

```
dscli> showvolgrp  IBM.2107-75FA120/V11
```

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 1

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output reports that are associated with the **showvolgrp** command.

Note: The volume group type determines the format of the LUN ID that is reported. The following examples demonstrate these differences.

Invoking the showvolgrp command where the volume group type is SCSI MAP 256

```
dscli> showvolgrp -lunmap IBM.2107-1300861/V2
```

The resulting output

Name	ID	Type	Vols
My_host_ system_ volumes	V2	SCSI- MAP 256	1000 1001 1002 1003 1004 1005 1006 1007

```
============================LUN Mapping==========================
```

vol	lun
1000	00
1001	01
1002	02
1003	03
1004	04
1005	05
1006	06
1007	07

Example 2

Invoking the showvolgrp command where the volume group type is SCSI Mask

```
dscli> showvolgrp -etccmap IBM.2107-1300861/V18
```

The resulting output

Name	ID	Type	Vols
myVG1	V18	SCSI- Mask	1000 1001 1002 1003 1004 1005 1006 1007

```
============================ETCC Mapping==========================
```

vol	lun	ETCCstatus
1000	40104000	Unmanaged
1001	40104001	Unmanaged
1002	40104002	Managed
1003	40104003	Managed
1004	40104004	Unmanaged
1005	40104005	Unmanaged
1006	40104006	Unmanaged
1007	40104007	Unmanaged

Report field definitions

Name
Indicates the name that you assigned for the designated volume group ID.

ID Indicates the volume group ID. The volume group identifier is a four-digit decimal number having no leading zeros, and prefixed by a *V*.

Type
Indicates the configured volume group type. Any one of the following volume group types can be queried: FICON/ESCON All | SCSI all | SCSI Mask | SCSI MAP 256 | os400 all | os400 Mask | ESSNet Copy Services

Note: os400 all and **os400 Mask** are sometimes referred to as **SCSI520 all** and **SCSI520 Mask**.

Vols
Indicates the complement of accessible volume numbers within the designated volume group.

vol (part of LUN mapping table)
Indicates the volume ID.

lun (part of LUN mapping table)
Indicates the LUN ID that is mapped to the designated volume ID. As noted in the examples, the LUN IDs can be different based on volume group type.

ETCCstatus
Indicates the Easy Tier Cooperative Caching status.

Disabled
ETCC has been disabled for the storage system.

Managed
Currently, ETCC is managing the volume.

Unmanaged
Currently, ETCC is not managing the volume.

Unmonitored
Currently, ETCC is not monitoring the volume.

Unsupported
The storage system does not support ETCC.

Unknown
The ETCC status is unknown.

Advanced operation commands

Advanced operation commands are used to further administer and tune storage.

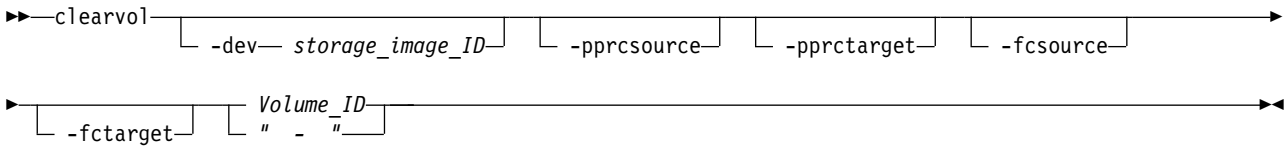
The following advanced operation commands are available:

clearvol
Clears Copy Services relationships for a base logical volume.

lsvolinit
Displays a list of fixed block volumes, either newly created or resized, that are still initializing in the ESS image. CKD volumes are not displayed.

clearvol

The **clearvol** command clears Copy Services relationships for a base logical volume.



Parameters

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-pprcsource
(Optional). This parameter is used with a base logical volume. It removes any remote mirror and copy relationships on the logical volume where the specified logical volume operates as a remote mirror and copy source.

-pprctarget
(Optional). This parameter is used with a base logical volume. It removes any remote mirror and copy relationships on the logical volume where the specified logical volume operates as a remote mirror and copy target.

-fcsource
(Optional). This parameter is used with a base logical volume. It removes any FlashCopy relationships on the logical volume where the specified logical volume operates as a FlashCopy source.

-fctarget
(Optional). This parameter is used with a base logical volume. It removes any FlashCopy relationships on the logical volume where the specified logical volume operates as a FlashCopy target.

Volume_ID | **-**
(Required). Specifies the volume ID where Copy Services relationships are to be cleared. This parameter accepts a fully qualified volume ID, which includes the storage image ID or a shortened version, if the **-dev** parameter is specified. The volume ID is a 32 bit number that can be represented as four hexadecimal digits in the form of *XYZZ* where:

X Specifies the address group, 0-F.

XY Specifies the logical subsystem number, 00 - FE.

ZZ Specifies the volume number, 00 - FF.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example: IBM.2107-75FA120/0001

Example

Invoking the clearvol command

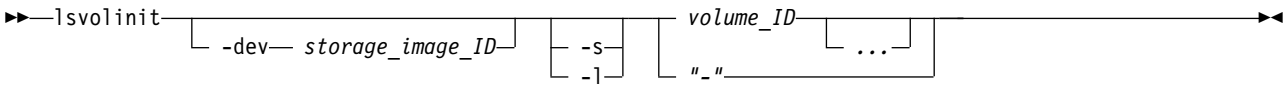
```
dscli> clearvol -dev IBM.2107-75FA120 0001
```

The resulting output

```
Volume 0001 successfully cleared.
```

lsvolinit

The **lsvolinit** command displays volumes that are being initialized using flash init. CKD volumes are not displayed.



Parameters

Note: If trackstoinit reaches zero, the initialization is complete and the volume is no longer listed with this command

-dev *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. Displays only the objects for the storage unit that is specified. The storage image ID is required if you do not specify a fully qualified volume ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-s
(Optional). Displays only the volume ID. You cannot use the **-s** and the **-l** parameters together.

-l
(Optional). Displays the default output. You cannot use the **-l** and the **-s** parameters together.

volume_ID ... | **-**
(Required). Specifies the IDs of the volumes that you want to query.

You must separate multiple volume IDs or ranges of volume IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lsvolinit** command using the **-l** parameter.

Invoking the lsvolinit command

```
dscli> lsvolinit -dev IBM.2107-75FA120 -l 0100-0101
```

The resulting output

ID	state	trackstoinit
100	Valid	128
101	Valid	1

Report field definitions

ID*
Indicates the volume ID.

State
Indicates the state of the initialization. One of the following values is displayed:

Valid
Indicates that the volume with the initialization in process is in a normal state and was queried successfully.

Validation Required
Indicates that the volume cannot be queried due to temporary microcode conditions. Issuing the query again, after a short interval, should solve the condition.

Volume Inaccessible
Indicates that the volume cannot be accessed (usually because it is in a fenced state) and the query failed.

Invalid
Indicates that a general internal error occurred during the processing of the query.

Note: If the state is anything other than Valid, then all other columns, except ID, are reported as " - ".

TracksToInit
Indicates the number of tracks that are not yet initialized. The maximum value that can be displayed is dependent on the volume size.

Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

Space-efficient storage commands

Space-effiicient storage commands create, modify, delete, and display extent space-efficient (ESE) and track space-efficient (TSE) repositories. These two types of repositories differ in many ways, but the main idea of a repository is identical.

A repository of type T and of capacity N guarantees that at least N bytes are available in the extent pool to provision the defined virtual T volumes. This configuration means that other types of non-T allocations will not consume all of the available storage in the extent pool. It also guarantees that no more than N bytes from the extent pool will be used to provision the defined virtual T volumes. Therefore, T allocations will not consume all of the available storage in the extent pool.

A repository also provides a clear and intuitive definition for the over-provisioned ratio, also known as the over-allocation or over-commitment ratio. The over-provisioned ratio is the sum of all the virtual T volume capacities divided by the repository capacity N. In other words, for any fixed repository capacity, the over-provisioned ratio varies only with virtual T volume creations and deletions. The ratio does not vary with any virtual non-T volume creations and deletions.

Although ESE volumes previously existed, the DS8000 system did not allow the creation of an ESE repository until Release 7.2. Without an ESE repository, no method to reserve storage to provision ESE volumes was available. Also, no method to prevent the ESE volume allocations from consuming the entire extent pool was available. Both problems were solved with the creation of an ESE repository, which is similar in effect to a TSE repository. However, the underlying mechanism to provide an ESE repository and the mechanism to provide a TSE repository are completely different.

Note: Track space-efficient (TSE) repositories are required to create TSE volumes. However, starting with DS8000 Release 8, TSE repositories can no longer be created.

For TSE repositories:

- All storage used to provision TSE volumes is physically allocated at repository-creation time.
- A TSE repository is required to create TSE volumes.
- The minimum TSE repository capacity is 16 GiB/Mod1.
- Changing the capacity size of an existing TSE repository is not supported.
- Changing the extent pool capacity does not affect the TSE repository capacity (capacity is fixed).
- All TSE volumes must be removed before attempting to remove the TSE repository.

For ESE repositories:

- All storage used to provision ESE volumes is physically allocated at volume-creation time.
- An ESE repository is not required to create ESE volumes.
- The minimum ESE repository capacity is 0 GiB/Mod1.
- Changing the capacity size of an existing TSE repository is supported.
- Changing the extent pool capacity does affect the ESE repository capacity (percentage is fixed).
- Any existing ESE volumes need not be removed before attempting to remove the ESE repository. The existing volumes remain unchanged after the repository is removed.

The following space-efficient storage commands are available:

chsestg
: Modifies the attributes of the specified repository in an extent pool.

lssestg
: Generates a report that displays the repository values for the entire storage image.

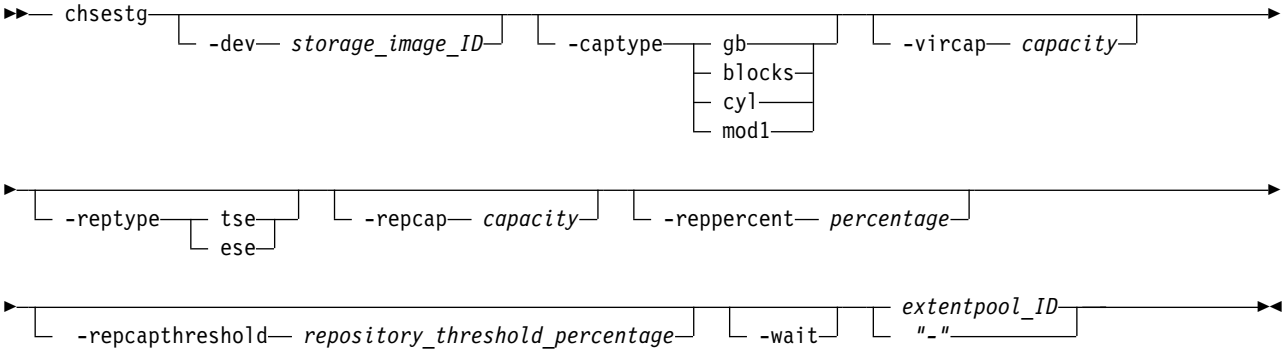
mksestg
: Creates the specified repository in an extent pool.

rmsestg
: Deletes the specified repository from an extent pool.

showsestg
: Generates a report that displays the detailed properties of the specified repository in an individual extent pool.

chsestg

The **chsestg** command changes the space-efficient storage attributes for an extent pool.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a value for manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the extent pool ID. It is also required if you do not set the *devid* variable in your profile or through the **setenv** command and the HMC is aware of more than one storage image. Using the *-dev* parameter temporarily overrides any defined value for *devid* for the current command.

-captype *gb* | *blocks* | *cyl* | *mod1*
(Optional) Specifies the type of capacity unit that is specified by the **-vircap** and **-repcap** parameters. *gb* is the default for fixed block storage, and *mod1* is the default for count key data (CKD). You can also specify the -captype gb parameter on CKD storage.

Select one of the following capacity unit types:

gb (Gibibyte)
Specifies the capacity unit type as gibibytes (GiB). 1 GiB = 2 097 152 blocks.

blocks Specifies the capacity unit type as blocks. This capacity unit type can be selected only for fixed block storage. One block = 512 Bytes.

cyl Specifies the capacity unit type as cylinders. This capacity unit type can be selected only for CKD storage. One cylinder = 15 tracks.

mod1 Specifies the capacity unit type in mod1 units. This capacity unit type can be selected only for CKD storage (mod1 = 1113 cylinders). This type is the default for CKD storage when the **-captype** parameter is not specified.

-vircap *capacity*
(Optional) Specifies the amount of virtual capacity that can be allocated to all space-efficient logical volumes, including ESE volumes. All capacities must be designated as whole numbers. The capacity units are specified by the **-captype** parameter.

If a -**vircap** parameter capacity is specified and it is larger than the existing vircap capacity, the specified capacity replaces the existing capacity. Otherwise, the specified capacity is ignored.

Notes:

1. Some releases of the DS8000 require that the ratio of virtual capacity to repository capacity to be at least 2:1.
2. Only one of the following parameters can be specified: **-repcapthreshold**, **–repcap**, **-reppercent**, and **-vircap**.

-reptype ese
(Optional) Specifies the type of repository that is affected by the **repcap**, **reppercent**, and **repcapthreshold** parameters..

Note: The default for the **reptype** parameter depends on the release level of your storage system. For example, starting with DS8000 Release 8.0 storage systems, a TSE repository can no longer be created and is no longer supported. For Release 8.0 and later, if you do not specify the **reptype** parameter, the default is ESE.

For pre-Release 8.0 storage systems, if you do not specify the **reptype** parameter, the default is TSE.

tse
Specifies a track space-efficient physical repository.

ese
Specifies an extent space-efficient physical repository. This value is the default.

Note: An ESE repository is not required to create ESE volumes. By creating an ESE repository, you specify both a minimum capacity reserved for ESE volumes and a maximum capacity allowed for ESE volumes. By default, no ESE repository allows the entire pool to be used by any requestor of that capacity whether it is for standard volumes, a TSE repository, virtual capacity (overhead), or ESE volumes. That is, the ESE volumes have 0% of the pool that is guaranteed, and 100% the extent limit of the pool is allowed.

-repcap *capacity*
(Optional) Specifies the amount of physical capacity for the specified repository that you want to provision the virtual capacity of the space-efficient logical volumes. The repository type is specified with the **-reptype** parameter. All capacities must be designated as whole numbers. The capacity units are specified by the -**captype** parameter. Only one of the following parameters can be specified in a single command: **-repcap**, **-reppercent**, **-vircap**.

The following rules apply to ESE repositories:

- An ESE repository is not required to create ESE volumes.
- Changing the capacity size of an existing ESE repository is supported. The actual repository capacity is equal to the specified capacity, rounded up to the nearest whole percentage of the pool capacity.
- The **chsestg** command cannot be used to create an ESE repository. Use the **mksestg** command.
- The **chsestg** command cannot be used to remove an ESE repository. Use the **rmsestg** command.
- Existing ESE volumes do not need to be removed before you attempt to remove the ESE repository. The existing volumes remain unchanged after the repository is removed.
- The repository can be resized to any capacity from the size of the capacity that is used to provision any existing ESE volumes to smaller than the available capacity. The virtual storage overhead capacity is not a part of the repository. Therefore, the largest repository capacity is equal to the total available capacity, less the overhead capacity, rounded up to the nearest whole percentage of the pool capacity.

-reppercent *percentage*
(Optional) Specifies the amount of physical capacity for the specified repository that you want to provision the virtual capacity of the space-efficient logical volumes. The value is specified as a percentage of the total virtual capacity. The repository type is specified with the **-reptype** parameter. Only one of the following parameters can be specified in a single command: **-repcap**, **-reppercent**, **-vircap**.

This parameter is provided as a convenience. It is the equivalent to using the **-repcap** parameter with a value determined by multiplying the total virtual capacity by the specified percentage and rounding up to the nearest whole number. However, this parameter is seldom used because the total virtual capacity is automatically increased as space-efficient volumes are created. While the virtual capacity can be set manually to a higher value, the total virtual capacity is equal to the sum of all capacities of space-efficient volumes.

-repcapthreshold *repository_threshold_percentage*
(Optional) Specifies the minimum threshold percentage of the physical repository capacity available. When the percentage of the currently available repository capacity is less than this minimum percentage, notifications are sent and the repository capacity status is reported as exceeded.

Notes:

1. Three thresholds for the repository generate notifications when their thresholds amounts are attained. Two of the three thresholds are set by the system and cannot be changed. They are set to 0% (full) and 15% (85% full). The third threshold is the user-defined threshold that is set here, and the repository capacity status is based on this threshold. When any of the three thresholds attain a threshold amount, a notification is sent for that particular threshold. No further notifications are sent until the repository capacity changes. If the repository capacity changes and remains above the threshold, another notification might be sent, but no more than one notification every 5 minutes. You must free capacity in the repository to stop the notifications. If the user-defined

threshold is equal to one of the other two fixed thresholds, only one notification is sent, at most once every 5 minutes, for the two equivalent thresholds.

2. To verify that your storage complex is set up to send notifications, use the **showsp** command. If it is not set up, use the **chsp** command to set up notifications.

-wait
(Optional) Specifies that the command will be delayed until after the space efficient storage is created, configured, and in a Normal state. If an error condition is detected while waiting, the command returns and reports an error. The **-wait** parameter can be specified if either **-vircap**, **-repcap**, or **-reppercent** is also specified.

extentpool_ID | –
(Required) Specifies the ID of the space-efficient storage extent pool that you want to change. This parameter accepts either a fully qualified extent pool ID or a shortened version if the **-dev** parameter is used. The shortened version is a four-digit decimal number with no leading zeros, prefixed with the letter *P*.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the chsestg command to modify space-efficient storage in an extent pool.

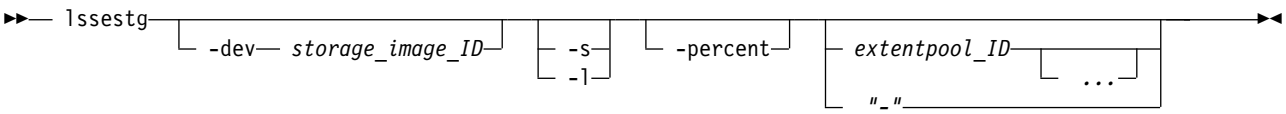
```
dscli> chsestg -dev IBM.2107-75FA120 –repcapthreshold 75 P2
```

The resulting output

```
The space-efficient storage for the extent pool P2
has been modified successfully.
```

lssestg

The **lssestg** command displays a list of the space-efficient repositories in the storage system.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a value for manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified extent pool ID. It is also required if you do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-s
(Optional) Displays only the extent pool IDs of the extent pools that contain space-efficient storage. You cannot use the **-s** and the **-l** parameters together.

-l
(Optional) Displays the default output plus the virtual and physical repository capacity that is allocated for space-efficient repositories. You cannot use the **-s** and the **-l** parameters together.

-percent
(Optional) Specifies that the repcapalloc and vircapalloc values be displayed in percentages rather than in gibibyte (GiB) or Mod1 units.

Note: In some versions of the DS CLI, the displayed percentages for the repcapalloc and vircapalloc values used one decimal place. However, because internal percentage calculations only use whole numbers, the decimal place was removed.

extentpool_ID ... | -
(Optional) Specifies the IDs of one or more extent pools that you want the system to display the space-efficient storage details for. A fully qualified extent pool ID is accepted, which consists of the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is a four-decimal digit number with no leading zeros, prefixed with the letter P.

To specify a range of extent pool IDs, separate the extent pool IDs with a hyphen.

You must separate multiple extent pool IDs or ranges of extent pool IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lssestg** command by using the **-l** parameter.

Invoking the lssestg command to display space-efficient repositories in a storage system.

```
dscli> lssestg -dev IBM.2107-75FA120 -l
```

The resulting output

```
dscli> lssestg -l
Date/Time: August 24, 2016 9:21:08 AM MST IBM
DS: IBM.2107-75YZ881
```

extpool	stgtype	datastate	configstate	repcapstatus
P5	ckd	Normal	Normal	below
P10	fb	Normal	Normal	full
P13	fb	Normal	Normal	-
P20	fb	Normal	Normal	below
P20	fb	Normal	Normal	below

repcap(GiB/ Mod1)	vircap	reptype	repcapalloc	vircapalloc	opratio
100.0	200.0	tse	10.0	150.0	1.5
0.0	0.0	ese	0.0	0.0	-
0.0	100.0	-	-	0.0	-
160.0	1600.0	tse	0.0	0.0	0.0
1040.4	1600.0	ese	0.0	500.0	0.5

Notes:

1. Starting with DS8000 Release 8.0, a TSE repository can no longer be created and is no longer supported. For more information, see the **reptype** parameter included in the **mksestg** command.
2. The extent pool P5 has a TSE repository, but no ESE repository, and an over-provisioned ratio of 1.5 to 1.
3. The extent pool P10 has an ESE repository that is created with 0 (zero) capacity, which prevents ESE volume creation in the pool. The `repcapstatus` is `full` because the pool does not allow any further ESE volume creation.
4. The extent pool P13 has the virtual capacity that is specified but a TSE or an ESE repository. ESE volume creation is still allowed. However, because an ESE repository does not exist, `vircapalloc` is always zero. See the `virconfigured` field of the **showextpool** command to see total space-efficient storage of the extent pool.
5. The extent pool P20 has both a TSE and an ESE repository. The value for the `vircap` column is the same for both repositories because the virtual capacity is defined for all space-efficient storage.

Report field definitions

Extpool*
Identifies the extent pool that you are querying.

Stgtype
Identifies the storage type. The value that is displayed is either fb (fixed block) or ckd (count key data).

Datastate
One of the following data states are displayed:

Normal
Indicates that the space-efficient storage state is normal and that none of the other data states apply.

Pinned
Indicates that none of the other data states apply and that logical tracks are present or at least identified in NVS or cache and cannot be de-staged for one reason or another.

Read only
Indicates that the logical volume is read only because one or more extents on the logical volume are on a rank in the read only data state.

Inaccessible
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the inaccessible data state.

Indeterminate data loss
Indicates that the following data states do not apply and that one of the following conditions occurred:

Data states that do not apply:
- Rank failed
- Rank repairing
- Rank repaired
- Global inaccessible
- Global lost data

Conditions - one of the following items occurred:
- Committed write data was lost before it was de-staged, and the track identifiers that are associated with the data are unknown.
- Data was lost that indicates that extents on the logical volume were active FlashCopy targets.

Rank failed
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the Failed data state. This data state moves to the Rank repairing state if the rank moves to the Rank repairing state through use of the repair array function.

Rank Repairing
Indicates that one or more extents that are associated with the logical volume are on ranks in the repairing data state.

Rank Repaired
Indicates that one or more extents that are associated with the logical volume are on ranks that were in the repairing state, but are not in the repairing state now.

Global inaccessible
Indicates that the global metadata that is associated with the logical volume configuration is inaccessible. Some of the data that is associated with the logical volume might be inaccurate.

Global lost data
Indicates that global metadata that is associated with the logical volume configuration was lost. As a result, some of the data that is associated with the logical volume might be inaccurate.

NVS data inaccessible
Indicates that active NVS data is inaccessible for one or more logical volumes of an LSS group. The logical volumes in the LSS group cannot be made accessible.

Extent fault
Indicates that none of the other states apply and a logical volume needs virtual space that is converted to real space, but the space was not available. So the subsequent writes fail until the space becomes available.

Configstate
One of the following configuration states is displayed:

Normal
Indicates that no space-efficient storage configuration operations are in progress.

Configuration pending
Indicates that an initial configuration for space-efficient storage is in the queue.

Configuration pending error
Indicates that the request for initial configuration for space-efficient storage to be in the queue did not complete successfully.

Configuring
Indicates that space-efficient storage is being configured for the first time.

Configuration error
Indicates that the initial configuration did not complete successfully. This state reflects an internal error condition and not an error in the request to create the space-efficient storage.

Corrective action: Use the**rmsestg** command to delete each track in the space-efficient storage that is listed with the configuration state of "configuration error".

Reconfiguration error
Indicates that the reconfiguration request did not complete successfully.

Migration error
Indicates that the dynamic volume relocation operation was ended during processing.

Configuration out-of-synch
Indicates that internal inconsistencies exist for the configuration state of the space-efficient storage.

Deconfiguring
Indicates that the space-efficient storage is being deleted.

Deconfiguration error
Indicates that a request to delete space-efficient storage did not complete successfully. This state reflects an internal error condition and not an error in the request to remove the volume. To correct this state, you must reissue the **rmsestg** command for the space-efficient storage that is listed with the configuration state of "deconfiguration error".

Degraded - Configuration Error
Indicates that some of the storage configuration process failed to complete successfully. Some is normal and it can continue to be used.

Degraded - Configuration Out of Synch
Indicates that internal inconsistencies exists for the configuration state of some of the storage. Some is normal and it can continue to be used.

Degraded - Configuration Pending
Indicates that the configuration operation for some of the storage is queued. Some is normal and it can continue to be used.

Degraded - Configuring
Indicates that some of the storage is in the process of configuring. Some is normal and it can continue to be used.

Degraded - Deconfiguration Error
Indicates that some of the deconfiguration process did not complete successfully. Some is normal and it can continue to be used.

Degraded Deconfiguring
Indicates that some of the storage is being deleted. Some is normal and it can continue to be used.

Degraded - Migration Error
Indicates that some of the space-efficient storage is being migrated and some is normal and it can continue to be used.

Degraded - Pending
Indicates that some of the configuration operation is queued. Some is normal and it can continue to be used.

Degraded - Reconfiguration Error
Indicates that some of the space-efficient storage is being reconfigured and some is normal and can continue to be used.

Unknown
Indicates that the configuration state of the space-efficient storage cannot be determined due to an internal error.

Partial - No Physical Space
Indicates that no physical space is defined. Defining physical space is not required and the state might be normal, with only virtual space defined.

Partial - No Virtual Space
Indicates that the physical space is defined, but the virtual space is not.

Reconfiguring
Indicates that the space-efficient storage is being reconfigured.

Degraded -Reconfiguring
Indicates that some of the space-efficient storage is being reconfigured and some is normal and that it can continue to be used.

Migrating
Indicates that space-efficient storage is being migrated.

Degraded - Migrating
Indicates that some of the space-efficient storage is being migrated and some is normal and that it can continue to be used.

Merging
Indicates that the volume is in the process of merging. For example, merging from one extent pool to a different extent pool.

Degraded - Merging
Indicates that some of the space-efficient storage is being merged and some is normal and that it can continue to be used.

Transposition Error
Indicates that an internal error condition occurred. This error can happen when a merge extent pool operation fails. To correct this state, use the **chextpool** command with the **-merge** parameter to redrive the original merge pool operation.

Degraded - Transposition Error
Indicates that an internal error condition occurred on some of the space-efficient storage and some is normal and that it can continue to be used.

Repcapstatus
Indicates the status of the repository capacity. One of the following three values is displayed:

- A dash (-) is displayed if the status is undefined or not applicable. For example, a dash is displayed if the repository does not exist.

below
The repository capacity available (repcap - repcapalloc), as a percentage of total repository capacity (repcap), is greater than the repository capacity threshold.

exceeded
The repository capacity available is less than the repository capacity threshold.

full
The repository capacity available is zero.

Repcap(GiB/Mod1)
Indicates the total physical repository capacity in the format of X.Y *where*, for fixed block volumes, X is in whole gibibytes (GiB) and Y represents tenths of a GiB, and is limited to a single digit (0 - 9). For CKD volumes, the Mod1 capacity (1 Mod1 = 1113 cylinders) is displayed, *where* X is a whole Mod1, and Y represents tenths of a Mod1.

Vircap
Indicates the total virtual capacity for all space-efficient volumes, including ESE and TSE, in the format of X.Y *where*, for fixed block volumes, X is in whole gibibytes (GiB) and Y represents tenths of a GiB, and is limited to a single digit (0 - 9). For CKD volumes, the Mod1 capacity (1 Mod1 = 1113 cylinders) is displayed, *where* X is a whole Mod1, and Y represents tenths of a Mod1.

Reptype
Indicates the type of repository, as follows:

tse
Indicates a track space-efficient (TSE) physical repository.

ese
Indicates an extent space-efficient physical repository.

- A dash (-) indicates that this information is unknown, not available, or not applicable.

Repcapalloc+
Indicates the allocated physical repository capacity of the space-efficient storage that is used from the available repository capacity as a result of writes to the space-efficient volume. Displayed in the format of X.Y *where*, for fixed block volumes, X is in whole gibibytes (GiB) and Y represents tenths of a GiB, and is limited to a single digit (0 - 9). For CKD volumes, the Mod1 capacity (1113 cylinders) is displayed, *where* X is a whole Mod1, and Y represents tenths of a Mod1.

If the **-percent** parameter is specified, the percent of total allocated physical repository is displayed in the format of a whole number with % (for example, 12%).

Vircapalloc+
Indicates the allocated virtual capacity of the specified space-efficient storage, that is to say, the amount of virtual capacity that is already allocated to logical capacity on the ESE or TSE volumes. If there is no specified repository, meaning the reptype value is a '-' (dash), this value is zero. This will be true, even if a particular pool contains ESE space-efficient volumes but does not contain an ESE repository.

Displayed in the format of X.Y *where*, for fixed block volumes, X is in whole gibibytes (GiB) and Y represents tenths of a GiB, and is limited to a single digit (0 - 9). For CKD volumes, the Mod1 capacity (1113 cylinders) is displayed, *where* X is a whole Mod1, and Y represents tenths of a Mod1.

If the **-percent** parameter is specified, the percent of total virtual capacity that is defined as space-efficient volumes is displayed in the format of a whole number with % (for example, 12%).

Opratio+
Over-provisioned ratio is the allocated virtual capacity that is divided by the size of the repository (vircapalloc / repcap). Displayed in the format of X.Y *where* Y is limited to a single digit (0 - 9).

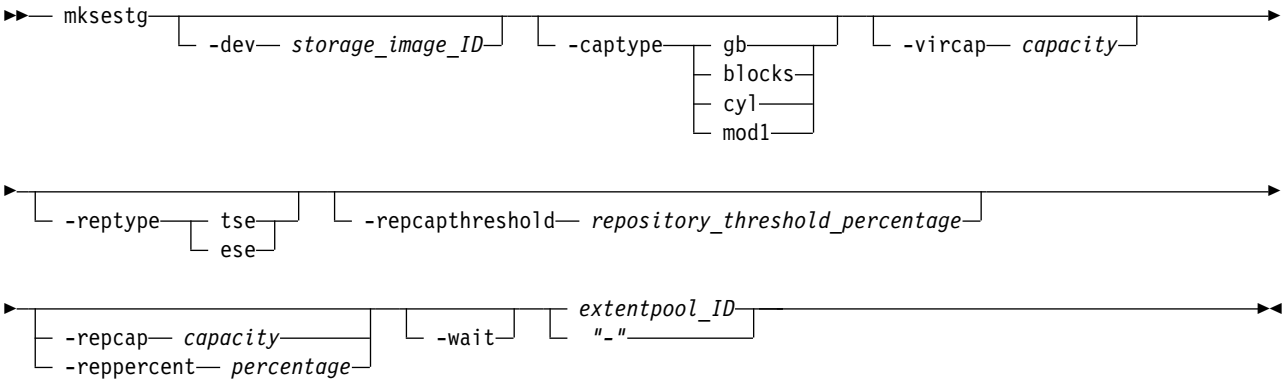
Key:

* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

mksestg

The **mksestg** command creates a space-efficient repository in an existing extent pool.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a value for manufacturer, machine type, and serial number. The *-dev* parameter temporarily overrides any defined value for *devid* for the current command. The storage image ID is required for the following situations:

- If you do not specify a fully qualified extent pool ID
- If you do not set the *devid* variable in your profile or through the **setenv** command

- If the HMC is aware of more than one storage image

-captype gb | blocks | cyl | mod1
(Optional) Specifies the unit type of capacity that is specified by the **-vircap** and the **–repcap** parameters. The following values are specified for each of these unit types:

gb The capacity unit is gibibytes (GiB). This type is the default for fixed block storage when the **-captype** parameter is not specified. 1 GiB = 2 097 152 blocks.

This unit type can also be used with count key data (CKD) storage.

blocks The capacity unit is blocks. This unit type is used only with fixed block storage. 1 block = 512 Bytes.

cyl The capacity unit is cylinders. This unit type is used only with count key data (CKD) storage. 1 cylinder = 15 tracks.

mod1 The capacity unit is mod1. This unit type is used only with CKD storage (1 mod1 = 1113 cylinders). This type is the default for CKD storage when the **-captype** parameter is not specified.

-vircap *capacity*
(Optional) Specifies the amount of virtual capacity that can be allocated to all space-efficient logical volumes, including TSE and ESE volumes. All capacities must be designated as whole numbers. The capacity units are specified by the **-captype** parameter.

Note: Some DS8000 models require the ratio of virtual capacity to repository capacity to be at least 2:1.

-reptype tse | ese
(Optional) Specifies the type of repository that is affected by the **repcap**, **reppercent**, and **repcapthreshold** parameters.

Note: The default for the **reptype** parameter depends on the release level of your storage system. For example, starting with DS8000 Release 8.0 storage systems, a TSE repository can no longer be created and is no longer supported. For Release 8.0 and later, if you do not specify the **reptype** parameter, the default is ESE.

For pre-Release 8.0 storage systems, if you do not specify the **reptype** parameter, the default is TSE.

tse
Specifies a track space-efficient physical repository.

ese
Specifies an extent space-efficient physical repository.

Note: An ESE repository is not required to create ESE volumes. By creating an ESE repository, you specify both a minimum capacity that is reserved for ESE volumes and a maximum capacity that is allowed for ESE volumes. By default, no ESE repository allows the entire pool to be used by any requester of that capacity whether it is for standard volumes, a TSE repository, virtual capacity (overhead), or ESE volumes. That is, the ESE volumes have 0% of the pool that is guaranteed, and the extent limit of the pool is allowed.

-repcapthreshold *repository_threshold_percentage*
(Optional) Specifies the minimum threshold percentage of the physical repository capacity that is available. When the percentage of the currently available repository capacity is less than this minimum percentage, notifications are sent and the repository capacity status is reported as exceeded. The default value is zero.

Notes:

1. Three thresholds for the repository generate notifications when their thresholds amounts are attained. Two of the three thresholds are set by the system and cannot be changed. They are set to 0% (full) and 15% (85% full). The third threshold is the user-defined threshold that is set here, and the repository capacity status is based on this threshold. When any of the three thresholds attain a threshold amount, a notification is sent for that particular threshold. No further notifications are sent until the repository capacity changes. If the repository capacity changes and remains above the threshold, another notification might be sent, but no more than one notification every 5 minutes. You must free capacity in the repository to stop the notifications. If the user-defined threshold is equal to one of the other two fixed thresholds, only one notification is sent, at most once every 5 minutes, for the two equivalent thresholds.
2. To verify that your storage complex is set up to send notifications, use the **showsp** command. If it is not set up, use the **chsp** command to set up notifications.

-repcap *capacity*

(Optional) Specifies the amount of physical capacity for the specified repository that you want to provision the virtual capacity of the space-efficient logical volumes. It is considered an error to create a repository that exists, but a single repository of each type is allowed in each extent pool. The repository type is specified with the **-reptype** parameter. All capacities must be designated as whole numbers. The capacity units are specified by the **-captype** parameter. You cannot specify the **-repcap** and **-reppercent** parameters at the same time.

Attention: TSE repositories can no longer be created starting with DS8900F Release 8.

An error occurs if you use this command to create a TSE/ESE repository when a TSE/ESE repository exists.

For TSE repositories:

- A TSE repository is required to create TSE volumes.
- All storage that is used to provision TSE volumes is physically allocated at the repository-creation time.
- The minimum repository capacity that can be created is as follows:
  - **gb** = 16 GiB
  - **blocks** = 33 554 432 blocks (equivalent to 16 GiB).
  - **cyl** = 17 808 cylinders.

For ESE repositories:

- An ESE repository is not required to create ESE volumes.
- All storage that is used to provision TSE volumes is physically allocated at the volume-creation time.
- The minimum repository capacity that can be created is 0 (zero) and prevents any ESE volume creation.
- Any existing ESE volumes become a part of the newly created ESE repository but remain otherwise unchanged. Any existing ESE volumes raise the minimum required repository capacity to be equal to capacity already allocated to provision those ESE volumes.

-reppercent *percentage*

(Optional) Specifies the amount of physical capacity for the specified repository that you want to provision the virtual capacity of the space efficient logical volumes. The physical capacity is specified as a percentage of the total virtual capacity. Creating a repository that exists is considered an error, but a single repository of each type is allowed in each extent pool. The repository type is specified with the **-reptype** parameter. You cannot specify the **-repcap** and **-reppercent** parameters at the same time. An error occurs if you use this command to create a TSE/ESE repository when a TSE/ESE repository exists.

This parameter is provided as a convenience. It is equivalent to using the **–repcap** parameter with a value determined by multiplying the total virtual capacity by the specified percentage and rounding

up to the nearest whole number. However, this parameter is seldom used because the total virtual capacity is automatically increased as space-efficient volumes are created. Therefore, while the virtual capacity can be manually set to a higher value, the total virtual capacity is equal to the sum of all the capacities of the space-efficient volumes.

Note: Some DS8000 models require the ratio of virtual capacity to repository capacity to be at least 2:1. On systems with this requirement, the effective maximum percentage is 50%. In many cases, a value of 20% of the virtual capacity is a good value. However, less than 20% of virtual capacity might significantly degrade performance.

-wait
(Optional) Specifies that the command will be delayed until after the space efficient storage is created, configured, and in a Normal state. If an error condition is detected while waiting, the command returns and reports an error.

extentpool_ID | **-**
(Required) Specifies the extent pool that is used to provision the extents that are used by this space efficient storage. For example, P111. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example 1

```
dscli> mksestg -dev IBM.2107-75FA120 -captype gb -vircap 32 -repcap 16 P101
```

The resulting output

```
The space-efficient storage for the extent pool P101 has been created
successfully.
```

Example 2

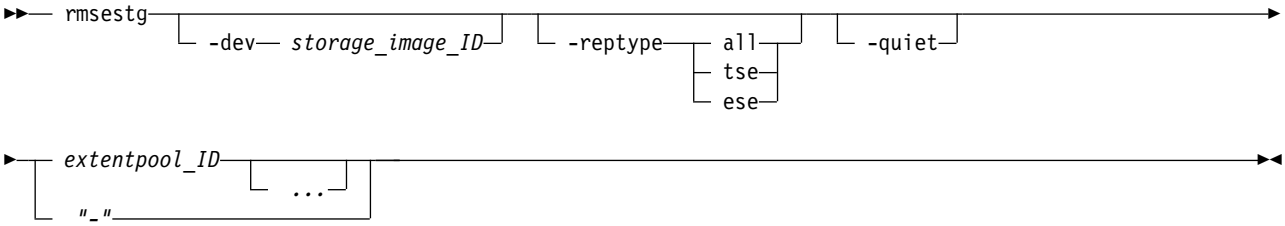
```
dscli> mksestg -vircap 1000 -repcap 500 -reptype tse P18
```

The resulting output

```
Track Space Efficient repositories are not supported.
```

rmsestg

The **rmsestg** command deletes the specified repositories in an extent pool.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a value for manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified ID for the extent pool ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-reptype all | **tse** | **ese**
(Optional) Specifies the type of repository to remove. If the type is not specified, the default is `all`.

Valid repository types are as follows:

all
Specifies the removal of all physical repositories, plus any overhead that is used to manage the virtual capacity. However, if no ESE volumes exist after this action but are later created, the virtual capacity overhead is automatically re-created.

tse
Specifies the removal of the physical TSE repository in an extent pool. Any ESE repository and virtual capacity overhead remain unchanged.

Notes:

1. Using this option is equivalent to using the **chsestg** command and setting the TSE repository capacity to zero (0).
2. A TSE repository is required to create TSE volumes. Starting with DS8000, Release 8, TSE repositories can no longer be created.
3. All TSE volumes must be removed before attempting to remove the TSE repository.

ese
Specifies the removal of the physical ESE repository in an extent pool. Any TSE repository and virtual capacity overhead remain unchanged.

Notes:

1. An ESE repository is not required to create ESE volumes.
2. Removing the ESE repository allows the entire pool to be used by any requestor of that capacity whether it is for standard volumes, a TSE repository, virtual capacity (overhead), or ESE volumes. That is, the ESE volumes have 0% of the pool guaranteed, and the extent limit of the pool is allowed)

.

-quiet
(Optional) Turns off the space-efficient storage removal confirmation prompt for this command.

extentpool_ID ... | -
(Required) Specifies the IDs of one or more extent pools that you want to delete the space-efficient storage from. A fully qualified extent pool ID is accepted, which consists of the storage image ID or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is a four-decimal digit number with no leading zeroes, prefixed with the letter P.

To specify a range of extent pool IDs, separate the extent pool IDs with a hyphen.

You must separate multiple extent pool IDs or ranges of extent pool IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input.

Note: You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the rmsestg command to create space-efficient storage in an extent pool.

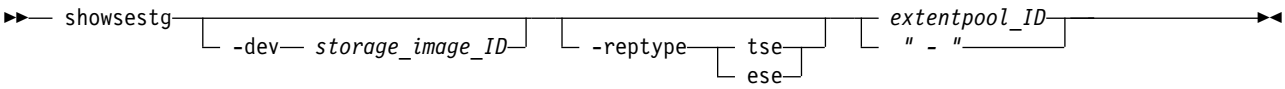
```
dscli> rmsestg -dev IBM.2107-75FA120 P2
```

The resulting output

```
Are you sure that you want to delete the space-efficient storage for
the extent pool P2?[Y/N]:
The space-efficient storage for the extent pool P2 has been
deleted successfully.
```

showsestg

The **showsestg** command displays a detailed properties report of the specified repository of an individual extent pool. Because an extent pool can have more than one repository, the **-reptype** parameter is used to specify the single repository to display.



Parameters

-dev *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of a value for manufacturer, machine type, and serial number. The storage image ID is required if you do not specify a fully qualified extent pool ID, do not set the *devid* variable in your profile or through the **setenv** command, and the HMC is aware of more than one storage image. Using the **-dev** parameter temporarily overrides any defined value for *devid* for the current command.

-reptype tse | ese
(Optional) Specifies the type of repository to display. All fields that begin with **rep** are associated with the specified repository type. These fields include **repcapstatus**, **%repcapthreshold**, **repcap**, **repcapalloc**, and **%repcapalloc**. If the type is not specified, the **-reptype** parameter defaults to tse.

Note: The default for the **reptype** parameter depends on the release level of your storage system. For example, starting with DS8000 Release 8.0 storage systems, a TSE repository can no longer be created and is no longer supported. For Release 8.0 and later, if you do not specify the **reptype** parameter, the default is ESE.

For pre-Release 8.0 storage systems, if you do not specify the **reptype** parameter, the default is TSE.

tse
Specifies a track space-efficient physical repository.

ese
Specifies an extent space-efficient physical repository.

extentpool_ID | –
(Required) Specifies the ID of the extent pool that you want to query for the space-efficient storage values. A fully qualified extent pool ID is accepted, which consists of the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is a four-decimal digit number with no leading zeros, prefixed with the letter P.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the showsestg command

```
dscli> showsestg –reptype ese p3
```

The resulting output

```
extpool               P3
stgtype               fb
datastate             Normal
configstate           Normal
repcapstatus          below
%repcapthreshold      0
repcap(GiB)           100.0
repcap(Mod1)          -
```

```
repcap(blocks)         209715200
repcap(cyl)            -
repcapalloc(GiB/Mod1)  24.0
%repcapalloc           24.0
vircap(GiB)            1000.0
vircap(Mod1)           -
vircap(blocks)         2097152000
vircap(cyl)            -
vircapalloc(GiB/Mod1)  48.0
%vircapalloc           4.8
overhead(GiB/Mod1)     2.0
reqrepcap(GiB/Mod1)    100.0
reqvircap(GiB/Mod1)    1000.0
reptype                ese
opratio                0.5
```

Report field definitions

Extpool
Indicates the extent pool that you are querying.

Stgtype
Indicates the storage type. The value that is displayed is either **fb** (fixed block) or **ckd** (count key data).

Datastate
One of the following data states is displayed:

Normal
Indicates that the space-efficient storage state is normal and that none of the other data states apply.

Pinned
Indicates that none of the other data states apply and that logical tracks are present or at least identified in NVS or cache that cannot be de-staged for one reason or another.

Read only
Indicates that the logical volume is read only because one or more extents on the logical volume are on a rank in the read-only data state.

Inaccessible
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the inaccessible data state.

Indeterminate data loss
Indicates that the following data states do not apply and that one of the following conditions has occurred:

Data states that do not apply:

- Rank failed
- Rank repairing
- Rank repaired
- Global inaccessible
- Global lost data

Conditions: one of the following has occurred:

- Committed write data was lost before it was de-staged, and the track identifiers that are associated with the data are unknown.
- Data has been lost that indicates that extents on the logical volume were active FlashCopy targets.

Rank failed
Indicates that one or more extents that are associated with the logical volume are on a rank that is in the Failed data state. This data state transitions to the Rank repairing state if the rank transitions to the Rank repairing state through use of the repair array function.

Rank Repairing
Indicates that one or more extents that are associated with the logical volume are on ranks in the repairing data state.

Rank Repaired
Indicates that one or more extents that are associated with the logical volume are on ranks that were in the repairing state, but are not in the repairing state now.

Global inaccessible
Indicates that the global metadata that is associated with the logical volume configuration is inaccessible. Some of the data that is associated with the logical volume might be inaccurate.

Global lost data
Indicates that global metadata that is associated with the logical volume configuration has been lost. As a result, some of the data that is associated with the logical volume might be inaccurate.

NVS data inaccessible
Indicates that active NVS data is inaccessible for one or more logical volumes of an LSS group. The logical volumes in the LSS group cannot be made accessible.

Extent fault
Indicates that none of the other states apply and a logical volume needs virtual space converted to real space, but the space was not available. So the subsequent writes fail until the space becomes available.

Configstate
One of the following configuration states is displayed:

Normal
Indicates that no space-efficient storage configuration operations is in progress.

Configuration pending
Indicates that an initial configuration for space-efficient storage is in the queue.

Configuration pending error
Indicates that the request for initial configuration for space-efficient storage to be in the queue did not complete successfully.

Configuring
Indicates that space-efficient storage is in the process of being configured for the first time.

Configuration error
Indicates that the initial configuration did not complete successfully. This state reflects an internal error condition and not an error in the request to create the space-efficient storage.

Corrective action: Use the **rmsestg** command to delete each track in the space-efficient storage that is listed with the configuration state of "configuration error".

Reconfiguration error
Indicates that the reconfiguration request did not complete successfully.

Migration error
Indicates that the dynamic volume relocation operation was ended during processing.

Configuration out-of-synch
Indicates internal inconsistencies for the configuration state of the space-efficient storage.

Deconfiguring
Indicates that the space-efficient storage is in the process of being deleted.

Deconfiguration error
Indicates that a request to delete space-efficient storage did not complete successfully. This state reflects an internal error condition and not an error in the request to remove the volume. To correct this state, you must reissue the **rmsestg** command for the space-efficient storage that is listed with the configuration state of "deconfiguration error".

Degraded - Configuration Error
Indicates that part of the storage configuration process failed to complete successfully.

Degraded - Configuration Out of Synch
Indicates internal inconsistencies for the configuration state of some of the storage. Some is normal, and it can continue to be used.

Degraded - Configuration Pending
Indicates that part of the configuration operation is queued. Some is normal and it can continue to be used.

Degraded - Configuring
Indicates that some of the storage is in the process of configuring. Some is normal and it can continue to be used.

Degraded - Deconfiguration Error
Indicates that part of the deconfiguration process did not complete successfully. Some is normal and it can continue to be used.

Degraded Deconfiguring
Indicates that some of the storage is in the process of being deleted. Some is normal and it can continue to be used.

Degraded - Migration Error
Indicates that some of the space-efficient storage is in the process of being migrated and some is normal and it can continue to be used.

Degraded - Pending
Indicates that the configuration operation is queued.

Degraded - Reconfiguration Error
Indicates that some of the space-efficient storage is in the process of being reconfigured and some is normal and can continue to be used.

Unknown
Indicates that the configuration state of the space-efficient storage cannot be determined due to an internal error.

Partial - No Physical Space
Indicates that there is no physical space defined. Defining physical space is not required and the state might be normal, with only virtual space defined.

Partial - No Virtual Space
Indicates that the physical space is defined, but the virtual space is not.

Reconfiguring
Indicates that the space-efficient storage is in the process of being reconfigured.

Degraded – Reconfiguring
Indicates that some of the space-efficient storage is in the process of being reconfigured and some is normal and that it can continue to be used.

Migrating
Indicates that space-efficient storage is in the process of being migrated.

Degraded - Migrating
Indicates that some of the space-efficient storage is in the process of being migrated and some is normal and that it can continue to be used.

Merging
Indicates that the volume is in the process of merging. For example, merging from one extent pool to a different extent pool.

Degraded - Merging
Indicates that some of the space-efficient storage is in the process of being merged and some is normal and that it can continue to be used.

Transposition Error
Indicates that an internal error condition has occurred. This error can happen when a merge extent pool operation fails. To correct this state, use the **chextpool** command with the **-merge** parameter to rerun the original merge pool operation.

Degraded - Transposition Error
Indicates that an internal error condition has occurred on some of the space-efficient storage and some is normal and that it can continue to be used.

Repcapstatus
Indicates the status of the repository capacity. One of the following three values is displayed:

- A dash (-) indicated that the status is undefined or not applicable. For example, a dash is displayed if the repository does not exist.

below
Indicates that the repository capacity available (repcap - repcapalloc), as a percentage of total repository capacity (repcap) is greater than the repository capacity threshold.

exceeded
Indicates that the repository capacity available is less than the repository capacity threshold.

full
Indicates that the repository capacity available is zero.

%repcapthreshold
Indicates the minimum threshold percentage of the physical repository capacity available. When the percentage of the currently available repository capacity is less than this minimum percentage, notifications are sent and the repository capacity status is reported as exceeded. The default value is zero.

Notes:

1. Three thresholds for the repository generate notifications when their thresholds amounts are attained. Two of the three thresholds are set by the system and cannot be changed. They are set to 0% (full) and 15% (85% full). The third threshold is the user-defined threshold that is set here, and the repository capacity status is based on this threshold. When any of the three thresholds have attained a threshold amount, a notification will be sent for that particular threshold. No further notifications will be sent until the repository capacity changes. If the repository capacity changes and remains above the threshold, another notification might be sent, but no more than one notification every five minutes. You must free capacity in the repository to stop the notifications. If the user-defined threshold is equal to one of the other two fixed thresholds, only one notification is sent, at most once every five minutes, for the two equivalent thresholds.
2. To verify that your storage complex is set up to send notifications, use the **showsp** command. If it is not set up, use the **chsp** command to set up notifications.

Repcap (GiB)
Indicates the total physical repository capacity in the format of X.Y, *where*, for fixed block volumes, X is in whole gibibytes (GiB) and Y represents tenths of a GiB, and is limited to a single digit (0 - 9).

Repcap (Mod1)
Indicates the total physical repository capacity for CKD in the format of X.Y, *where* X is in whole

Mod1 units (1113 cylinders) and Y represents tenths of a Mod1 unit, and is limited to a single digit (0 - 9). A value is displayed if the storage is for CKD, otherwise, a " - " value is displayed if the storage is for fixed block.

Repcap (blocks)
Indicates the total physical repository capacity in blocks. A value is displayed if the storage is for fixed block, otherwise, a " - " value is displayed if the storage is for CKD.

Repcap (cyl)
Indicates the total physical repository capacity in cylinders. A value is displayed if the storage is for CKD, otherwise, a " - " value is displayed if the storage is for fixed block.

Repcapalloc (GiB/Mod1)
Indicates the allocated physical repository capacity of the track space-efficient storage from the available repository capacity as a result of writes to the track space-efficient volumes. This value is displayed in the format of X.Y, *where*, for fixed block volumes, X is in whole gibibytes (GiB) and Y represents tenths of a GiB, and is limited to a single digit (0 - 9). For CKD volumes, X is in whole Mod1 units (1113 cylinders) and Y represents tenths of a Mod1 unit, and is limited to a single digit (0 - 9).

%repcapalloc
Indicates the allocated virtual capacity of the track space-efficient storage; that is, the amount of virtual capacity that is already defined as a percentage.

Note: In some versions of the DS CLI, the displayed percentage for the repcapalloc value used one decimal place. However, because internal percentage calculations only use whole numbers, the decimal place has been removed.

Vircap (GiB)
Indicates the total virtual capacity for all space-efficient volumes, including ESE and TSE. The format is X.Y, *where* X is in whole gibibytes (1 GiB) and Y represents tenths of a GiB, and is limited to a single digit (0 - 9).

Vircap (Mod1)
Indicates the total virtual capacity in the format of X.Y, where X is in whole Mod1 units (1113 cylinders) and Y represents tenths of a Mod1 unit, and is limited to a single digit (0 - 9). A value is displayed if the storage is for CKD, otherwise, a " - " value is displayed if the storage is for fixed block.

Vircap (blocks)
Indicates the total virtual capacity in blocks. A value is displayed if the storage is for fixed block, otherwise, a " - " value is displayed if the storage is for CKD.

Vircap (cyl)
Indicates the total virtual capacity in cylinders. A value is displayed if the storage is for CKD, otherwise, a " - " value is displayed if the storage is for fixed block.

Vircapalloc (GiB/Mod1)
Indicates the allocated virtual capacity of the specified space-efficient storage, that is to say, the amount of virtual capacity that is allocated to logical capacity on the ESE or TSE volumes. If there is no specified repository, meaning the reptype value is a '-' (dash), this value is zero. This will be true, even if a particular pool contains ESE space-efficient volumes but does not contain an ESE repository.

Displayed in the format of X.Y, *where*, for fixed block volumes, X is in whole gibibytes (GiB) and Y represents tenths of a GiB, and is limited to a single digit (0 - 9). For CKD volumes, X is in whole Mod1 units (1113 cylinders) and Y represents tenths of a Mod1 unit, and is limited to a single digit (0 - 9).

%vircapalloc
Indicates the space-efficient storage allocated virtual capacity (vircapalloc) as a percentage.

Note: In some versions of the DS CLI, the displayed percentage for the vircapalloc value used one decimal place. However, because internal percentage calculations only use whole numbers, the decimal place has been removed.

Overhead (GiB/Mod1)
Indicates the amount of physical space incurred to implement space-efficient storage.

reqrepcap(GiB/Mod1)
Indicates the total physical repository capacity that was requested in the format of X.Y, *where*, for fixed block volumes, X is in whole gibibytes (GiB) and Y represents tenths of a GiB, and is limited to a single digit (0 - 9). For CKD volumes, X is in whole Mod1 units (1113 cylinders) and Y represents tenths of a Mod1 unit, and is limited to a single digit (0 - 9).

reqvircap (GiB/Mod1)
Indicates the total virtual repository capacity requested in the format of X.Y, where X is in whole gibibytes (GiBs) and Y represents tenths of a GiB, and is limited to a single digit (0 - 9).

Reptype
(Optional) Indicates the type of repository, as follows:

tse
Indicates a track space-efficient physical repository.

ese
Indicates an extent space-efficient physical repository. This is the default value.

- A dash (-) indicates that this information is unknown, not available, or not applicable.

Opratio
Over-provisioned ratio is the allocated virtual capacity divided by the size of the repository (vircapalloc / repcap). Displayed in the format of X.Y *where* Y is limited to a single digit (0 - 9).

Storage image configuration commands

Commands are provided for information about configuring a storage image.

The following storage image configuration commands are available:

chsi Primarily enables or disables the ESSNet user interface that issues Copy Services operations. These processes are for the storage image, changes the description and name that you assigned to the storage image, or changes a System i serial number.

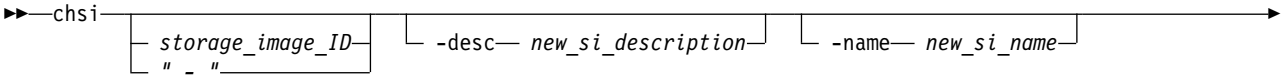
diagsi An administrative utility command that a user with administrator or service operator authority can use for non-routine diagnostic actions.

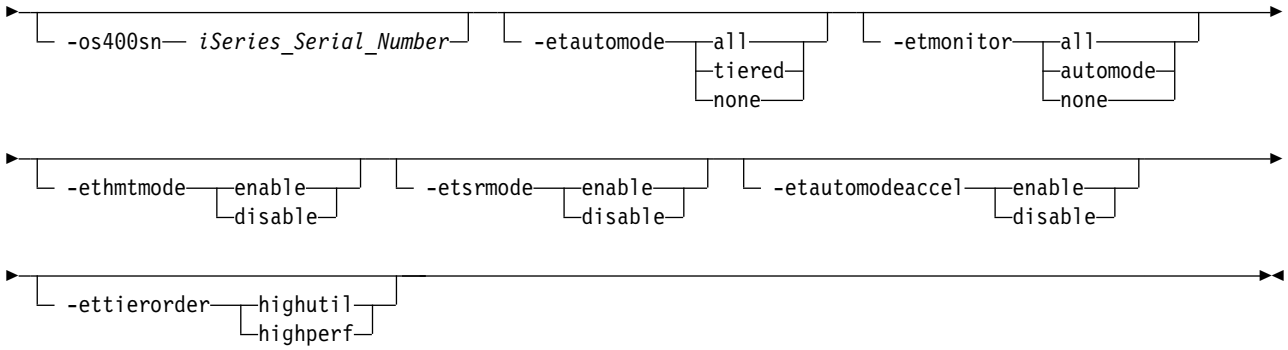
lssi Displays a list of storage images in a storage complex. These commands require that you use the storage image WWNN, which is displayed for each storage image when you use the **lssi** command.

showsi Displays the detailed properties of a storage image. In addition, the storage image WWNN is displayed for the specified storage image. The storage image WWNN is needed when you use the **lsavailpprcport** and **mkpprcpath** commands.

chsi

The **chsi** command modifies a storage image. You can use it to set characteristics such as online or offline state, name, and description. You can also use this command to select the allocation order that Easy Tier uses to determine which storage tiers are chosen when capacity is allocated in pools.





Parameters

storage_image_ID | **-**
(Optional) Specifies the storage image ID, which consists of the values for manufacturer, machine type, and serial number.

-desc *new_si_description*
(Optional) Specifies the description that you assign to the storage image. The description is limited to 256 bytes or 128 double-byte characters.

-name *new_si_name*
(Optional) Specifies the name that you assign to the storage image. The storage image name is limited to 32 byte or 16 double-byte characters.

-os400sn *iSeries_Serial_Number*
Specifies the new iSeries serial number.

The serial number consists of three hexadecimal characters. It uniquely identifies LUNs within a customer storage complex. It is appended to the unit serial number that is returned by a SCSI inquiry command that is directed to each LUN.

Notes:

1. You must restart both storage images after you process this DS CLI command to assign a new serial number.
2. The iSeries serial number is only required when you have multiple DS**** storage systems with the last 3 digits of the machine serial number that overlap.

-etautomode all | tiered | none
(Optional) Specifies the automatic mode of the IBM Easy Tier LIC feature. When the Easy Tier LIC feature is active, the manual mode is always active on all volumes, and the specified value controls only the automatic part of the feature.

all The Easy Tier automatic mode is active on all pools.

tiered The Easy Tier automatic mode is active only on pools with multiple tiers.

none The Easy Tier automatic mode is not active on any pools.

-etmonitor all | automode | none
(Optional) Specifies whether volumes are monitored by the IBM Easy Tier LIC feature. When the Easy Tier LIC feature is active, the volumes are monitored for the Easy Tier automatic mode.

all Easy Tier monitors all of the volumes on the DS8000 system, regardless of the state of the Easy Tier LIC feature. This capability demonstrates the potential benefits of Easy Tier. When this value is selected, all values of the **-etautomode** parameter are accepted, if the **IBM Easy Tier** LIC feature is active.

automode
Easy Tier monitors only those volumes that are managed by the Easy Tier automatic mode as specified by the **-etautomode** parameter.

none Easy Tier monitors none of the volumes on the DS8000 system. When this value is selected, all values of the **-etautomode** parameter are ignored, and the Easy Tier automatic mode is not active on any pools.

-etccmode enable | disable
(Optional) Easy Tier cooperative caching mode. The Easy Tier cooperative caching function is also known as Easy Tier Server. Only volumes that are monitored by Easy Tier can participate in Easy Tier cooperative caching.

Note: Starting with DS8000 Release 8, the **etccmode** parameter is no longer supported.

enable
Specifies that all volumes are allowed to participate in Easy Tier cooperative caching.

disable
Specifies that no volumes are allowed to participate in Easy Tier cooperative caching.

-ethmtmode enable | disable
(Optional) Specifies whether the IBM Easy Tier Heat Map Transfer feature is enabled or disabled on this storage facility image.

-etsrmode enable | disable
(Optional) Specifies whether space is reserved for Easy Tier data migration. The following values are included:

enable
Specifies that Easy Tier space reservation control is enabled. Enable is the default, which specifies that the required space for Easy Tier to do storage optimization actions is pre-reserved in each pool. The reserved space is counted as the pool overhead, and usually it is only a small amount of the space.

disable
Specifies that Easy Tier space reservation control is disabled. This value specifies that space is not pre-reserved for Easy Tier to do storage optimization actions. Ensure that you leave enough space in the storage pool for Easy Tier processing. If you select disable and the pool does not have enough free space, Easy Tier stops running storage optimization actions.

- Specifies that a dash (-) displays if the Easy Tier space reservation control is not supported on this storage system.

ETAutoModeAccel
Specifies the setting (enable or disable) of the Easy Tier Automatic Mode Acceleration control. When set to enable, the maximum amount of data that Easy Tier can automatically migrate in a given amount of time is doubled. This occurrence might cause your existing workload to experience increased latency and decreased bandwidth while migration occurs.

Note: Set this value to disable after the data that Easy Tier migrates is redistributed.

ETTierOrder
(Optional) Specifies the order in which drive classes are chosen when capacity is allocated in pools.

A dash (-) displays if the storage system does not support the Easy Tier allocation method.

Table 13. The table includes the tiers for the **highutil** *and* **highperf** *options according to storage type.*
Note: The default for hybrid storage systems is **highutil** and the default for All Flash storage systems is **highperf**.

highutil		highperf	
All Flash	**Hybrid**	**All Flash**	**Hybrid**
FlashTier1 (1)	Enterprise (2)	FlashTier0 (3)	FlashTier0
FlashTier2 (4)	FlashTier1	FlashTier1	FlashTier1
FlashTier0	FlashTier2	FlashTier2	FlashTier2
	Nearline (5)		Enterprise
	FlashTier0		Nearline

Notes:

1. The first tier of high capacity flash drives.
2. The tier that consists of drives with speeds of 10K RPM, 15K RPM, or a mixtures of 10K RPM and 15K RPM speeds.
3. The tier of high performance flash drives.
4. The second tier of high capacity flash drives.
5. The tier that consists of high-volume disks that are either SATA or SAS nearline drives.

Examples

Invoking the chsi command

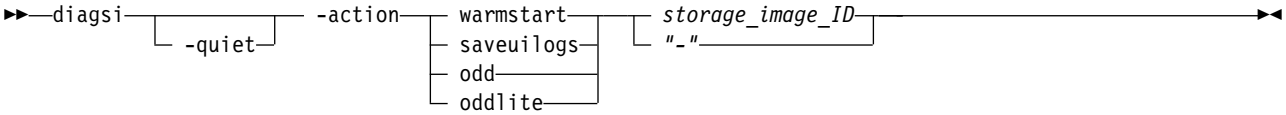
```
dscli> chsi -essnetcs n IBM.2107–75FA120
```

The resulting output

```
Storage image IBM.2107–75FA120 successfully modified.
```

diagsi

The **diagsi** command is an administrative utility command that a user with administrator or service operator authority can use for nonroutine diagnostic actions.



Parameters

Note: Only users with administrator or service authority are authorized to use this command.

-action warmstart | saveuilogs | odd | oddlite
(Required) Specifies the administrative action to be completed.

warmstart
The **-action** *warmstart* parameter initiates a warmstart on the storage image, which causes the storage image to collect microcode data that is useful in diagnosing problems. This action is restricted to the following usage rules:

- This action must be used only under the direction of your technical support representative.
- You must be in interactive mode to enter this action. You cannot enter this action while in single shot mode or from a script.
- Five minutes must pass before you can reissue the **-action** *warmstart* parameter.

- If you enter the **-action** *warmstart* parameter more than 10 times during a 24-hour period, the warmstart does not collect the microcode diagnostic data.

saveuilogs
The **-action** *saveuilogs* parameter offloads DS CLI logs to the Hardware Management Console and saves DS Network Interface, DS CLI, and DS8000 Storage Management GUI logs. All of the user interface log files can then be retrieved with a PE package.

odd
The **-action** *odd* parameter initiates an "on-demand dump" request.

oddlite
The **-action** *oddlite* parameter initiates a lightweight "on-demand dump" request.

-quiet
(Optional) Turns off the diagnostic control confirmation prompt for this command.

storage_image_ID | **-**
(Required) Specifies the fully qualified storage image ID. The storage image ID consists of manufacturer, machine type, and serial number.

If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) if you are using the DS CLI interactive mode.

Example

Invoking the diagsi command

```
dscli> diagsi -action oddlite IBM.2107-1300861
```

The resulting output

```
Are you sure you want to perform diagnostic control oddlite? [y/n]:y
CMUC00475I The diagnostic control action oddlite was submitted successfully.
```

Invoking the diagsi command

```
dscli> diagsi -action saveuilogs IBM.2107–68FA121
```

The resulting output

```
Are you sure you want to perform diagnostic control saveuilogs? [y/n]:y
CMUC00443I The diagnostic control action saveuilogs was submitted successfully.
The action will complete within 5 minutes.
```

lsserver

The **lsserver** command displays all servers in a storage complex or a list of specified servers and it also displays the status information for each server in the list.



Parameters

-s (Optional) Displays only the server ID. You cannot use the **-l** and the **-s** parameters together.

-l (Optional) Displays the default output and the state of the servers. You cannot use the **-l** and the **-s** parameters together.

Server_ID ... | **-**
(Optional) Displays the server information for the specified server IDs. This parameter accepts a fully

qualified server ID, which includes the storage image ID or a shortened version without the storage image ID. The shortened version is a two-digit decimal number with no leading zeros.

For DS8000,

- Example: IBM.2107-13AAV3A/0
- Example: IBM.2107-13AAV3A/1

To specify a range of server IDs, separate the server IDs with a hyphen.

You must separate multiple server IDs or ranges of server IDs with a blank space between each ID or range of IDs.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lsserver** command by using the **-l** parameter.

Invoking the lsserver command

```
dscli> lsserver -l
```

The resulting output

ID	Image ID	Image Name	Power Control SFI	State	LIC Version	OS Version	Bundle Version
IBM.2107 -75FA120 /00	1	SF13003 20ESS01	0	Online	5.3.0.832	-	-
IBM.2107 -75FA120 /01	1	SF13003 20ESS11	0	Online	5.3.0.832	-	-

Report field definitions

ID*
Indicates the unique identifier of the server. This value includes the storage image ID and the server ID.

Image ID
Indicates the image ID for the designated storage server.

Image Name
Indicates the image name for the designated storage server.

Power Control SFI
Indicates the storage server power control SFI.

State+
Indicates the current state of the designated server.

LIC Version
Indicates the LIC version for the designated storage server.

OS Version
Indicates the operating system version for the designated server. For a DS8000, this field always reports a dash.

Bundle Version
Indicates the bundle version for the designated storage server. For a DS8000, this field always reports a dash.

Key:

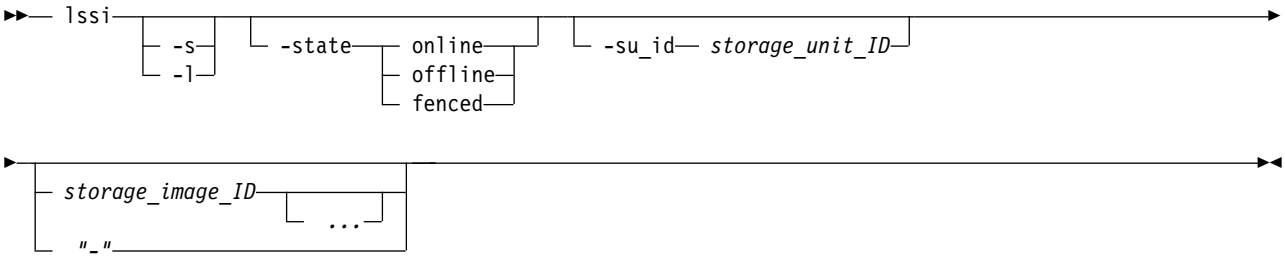
* Displayed when the **-s** parameter is specified.

+ Displayed only when the **-l** parameter is specified.

lssi

The **lssi** command displays a list of storage images in a storage complex.

You can use this command to look at the status of each storage image in the list. The storage image worldwide node name (WWNN) is displayed when this command is used. You must use the storage image WWNN with the **lsavailpprcport** and **mkpprcpath** commands.



Parameters

-s
(Optional) Displays only the storage image IDs. You cannot use the **-l** and the **-s** parameters together.

-l
(Optional) Displays the default output, ESSNet, volume group, and storage image description. You cannot use the **-l** and the **-s** parameters together.

-state online | offline | fenced
(Optional) Displays only the storage images in the specified state.

-su_id *storage_unit_ID . . .*
(Optional) Displays the storage images that are associated with the specified storage system. A storage system ID consists of manufacturer, machine type, and serial number.

storage_image_ID ... | **-**
(Optional) Accepts fully qualified storage image IDs. A storage image ID consists of manufacturer, machine type, and serial number. You must separate multiple IDs with a space between each ID.

Note: You cannot specify ID ranges.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **lssi** command by using the **-l** parameter. The input values differ when you use the **su_ID** and **storage_image_ID** parameters.

Invoking the lssi command

```
dscli> lssi -l
```

The resulting output

Name	ID	Storage System	Model	WWNN
DS 1	IBM.2107-75FA120	IBM.2107-75FA120	921	3007ACF30 12399E0

State	ESSNet	Volume Group	Desc
Online	Enabled	-	This is my DS storage Image

Report field definitions

Name
Indicates the name that you assigned to the storage system.

ID*
Indicates the storage image ID that consists of the manufacture, machine type, and serial number.

Storage System
Indicates the storage system ID that consists of the manufacture, machine type, and serial number.

Model
Indicates the model number that is associated with the storage system.

WWNN
Indicates the worldwide node name that is assigned to the storage system.

State
Indicates the status of the storage system. One of the following values are displayed:

Online
Indicates that the storage system is available to process all functions.

Offline
Indicates that the storage system is offline and not capable of processing any functions.

Resuming
Indicates that the storage system is preparing to come online.

Quiescing
Indicates that the storage system is preparing to go offline.

Quiesce Exception
Indicates that the storage system is in the quiesce exception state.

Forced Quiescing
Indicates that the storage system is preparing for a force offline operation.

Fenced
Indicates that the storage system failed and is offline.

ESSNet+
Indicates that the storage-complex ESSNet user interface can start Copy Services operations to this storage image. Enabled and Disabled are the values that are displayed in this field.

Volume Group+
Indicates the ESSNet Copy Services Volume Group ID or displays a " - " in this field.

If ESSNet Copy Services operations are enabled, the value that is displayed in this field specifies the ESSNet Copy Services type volume group. This volume group contains the logical volumes that can be controlled by Copy Services operations that are initiated through the ESSNet.

If ESSNet Copy Services operations are enabled and the ESSNet Copy Services Volume Group ID is not specified (represented by the " - " value in this field), all logical volumes are eligible to be controlled by Copy Services operations that are initiated through the ESSNet.

Desc+
Indicates the value that is assigned as a description for the storage system.

Key:

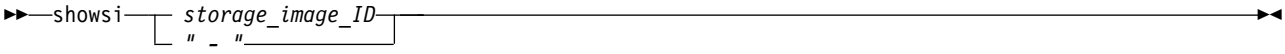
* Displayed when the **-s** parameter is specified.

\+ Displayed only when the **-l** parameter is specified.

showsi

The **showsi** command displays detailed properties of a storage image.

The storage image worldwide node name (WWNN) is displayed when this command is used. You must use the storage image WWNN with the **lsavailpprcport** and **mkpprcpath** commands.



Parameters

storage_image_ID | **-**
(Required) Specifies the storage image ID. A storage image ID consists of a manufacturer, machine type, and serial number.

If you use the dash (-), the specified value is read from standard input. However, you cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following tables represent the headers that are displayed on the output report that is associated with the **showsi** command.

Invoking the showsi command

```
dscli> showsi IBM.2107-75FA120
```

The resulting output

```
Name              sherwood
Desc              -
ID                IBM.2107-75NA901
Storage Unit      IBM.2107-75NA900
Model             996
WWNN              5005076408FEC795
Signature         9719-5d08-1865-1091
State             Online
ESSNet            Enabled
Volume Group      V0
os400Serial       795
NVS Memory        127.5 GB
Cache Memory      1867.3 GB
Processor Memory  1912.9 GB
MTS               IBM.5631-75NA910
numegsupported    16
ETAutoMode        all
ETMonitor         all
IOPMmode          Disabled
ETCCMode          -
ETHMTMode         Enabled
ETSRMode          Enabled
ETTierOrder       High utilization
| ETAutoModeAccel  Enabled
```

Report field definitions

Name
Specifies the name that you assigned to the storage system.

Desc
Specifies the value that you assigned as a description for the storage system.

ID Specifies the storage image ID that consists of the manufacture, machine type, and serial number.

Storage Unit
Specifies the storage system ID that consists of the manufacture, machine type, and serial number.

Model
Specifies the model number that is associated with the storage system.

WWNN
Specifies the worldwide node name that is assigned to the storage system.

Signature
Specifies the machine signature that is represented by 12 hexadecimal digits in the format *xxxx-xxxx-xxxx*.

State
Specifies the status of the storage system. One of the following values are displayed:

Online
Indicates that the storage system is available to process all functions.

Offline
Indicates that the storage system is not capable of processing any functions.

Resuming
Indicates that the storage system is starting to come online.

Quiescing
Indicates that the storage system is starting to go offline.

Quiesce Exception
Indicates that the storage system is in the quiesce exception state.

Forced Quiescing
Indicates that the storage system is starting a force offline operation.

Fenced
Indicates that the storage system has failed and is offline.

Discovery
Indicates that the storage system is determining which physical configurations are available and updates itself when it discovers new hardware.

ESSNet
Specifies that the storage-complex ESSNet user interface can start Copy Services operations to this storage image. Enabled or Disabled are the values that are displayed in this field.

Volume Group
Specifies the ESSNet Copy Services Volume Group ID or displays a " - " in this field.

If ESSNet Copy Services operations are enabled, the value that is displayed in this field specifies the ESSNet Copy Services type volume group. This volume group contains the logical volumes that can be controlled by Copy Services operations that are initiated through the ESSNet.

If ESSNet Copy Services operations are enabled and the ESSNet Copy Services Volume Group ID is not specified (represented by the " - " value in this field), all logical volumes are eligible to be controlled by Copy Services operations that are initiated through the ESSNet.

NVS Memory
Specifies the amount in decimal gigabytes (GB) of nonvolatile storage (NVS) memory that is configured on the storage system. Example 4.0 GB.

Cache Memory
Specifies the amount in decimal gigabytes (GB) of cache memory that is configured on the storage system. This memory equals the processor memory less the memory that is required for the operating system (varies with Licensed Internal Code level), NVS, and the tables that are required to manage the cache memory.

Processor Memory
Specifies the amount in decimal gigabytes (GB) of processor memory that is configured on the storage system. Processor memory equals the installed memory less the reserved memory requirements, including firmware memory, and varies with the DS8000 model, the DS8000 configuration, and the Licensed Internal Code level.

MTS
Specifies the order type of the storage system. The order type varies according to the warranty periods that are associated with the storage system.

numegsupported
Specifies the number of key groups that are supported.

ETAutoMode
Indicates the automatic mode of the IBM Easy Tier LIC feature when the Easy Tier feature is active. The manual mode is always active if the Easy Tier LIC feature is active.

all The Easy Tier automatic mode is active on all pools.

tiered The Easy Tier automatic mode is active only on pools with multiple tiers.

none The Easy Tier automatic mode is not active on any pools.

- The Easy Tier feature, both automatic and manual modes, is not supported.

ETMonitor
Specifies whether volumes are monitored by the IBM Easy Tier LIC feature.

all Easy Tier monitors all of the volumes on the system, regardless of the state of the Easy Tier LIC feature.

automode
Easy Tier monitors only those volumes that are managed by the Easy Tier automatic mode as specified by the **-etautomode** parameter.

none Easy Tier monitors none of the volumes on the system.

- The Easy Tier feature is not supported.

ETCCMode
Easy Tier cooperative caching. The Easy Tier cooperative caching function is also known as Easy Tier Server.

Note: Starting with Release 8, the **etccmode** parameter is no longer supported.

Enabled
Specifies that all volumes are allowed to participate in Easy Tier cooperative caching. Only volumes that are monitored by Easy Tier can participate in Easy Tier cooperative caching.

Disabled
Specifies that no volumes are allowed to participate in Easy Tier cooperative caching.

– Specifies that the DS8000 system does not support Easy Tier cooperative caching.

ETHMTMode
Specifies whether the IBM Easy Tier heat map transfer feature is enabled on this storage facility image.

Enabled
Specifies that the Easy Tier heat map transfer is enabled.

Disabled
Specifies that the Easy Tier heat map transfer feature is not enabled.

– Specifies that the storage system does not support the Easy Tier heat map transfer feature.

ETSRMode
Specifies whether space is reserved for Easy Tier data migration.

Enabled
Specifies that Easy Tier space reservation control is enabled. Enable is the default, which specifies that the required space for Easy Tier to do storage optimization actions is pre-reserved in each pool. The reserved space is counted as the pool overhead, and usually it is only a small amount of the space.

Disabled
Specifies that Easy Tier space reservation control is disabled. This value specifies that space is not pre-reserved for Easy Tier to do storage optimization actions. Ensure that you leave enough space in the storage pool for Easy Tier processing. If you select disable and the pool does not have enough free space, Easy Tier stops running storage optimization actions.

- Specifies that a dash (-) displays if the storage system does not support the Easy Tier space reservation function.

ETAutoModeAccel
Specifies the setting (enable or disable) of the Easy Tier Automatic Mode Acceleration control. When set to enabled, the maximum amount of data that Easy Tier can automatically migrate in a given amount of time is doubled. This occurrence might cause your existing workload to experience increased latency and decreased bandwidth while migration occurs.

Note: Set this value to disable after data that Easy Tier migrates is redistributed.

ETTierOrder
Specifies the order in which drive classes are chosen when capacity is allocated in pools.

A dash (-) displays if the storage system does not support the Easy Tier allocation method.

Table 14. The table includes the tiers for the **highutil** *and* **highperf** *options according to storage type.*
Note: The default for hybrid storage systems is **highutil** and the default for All Flash storage systems is **highperf**.

highutil		highperf	
All Flash	**Hybrid**	**All Flash**	**Hybrid**
FlashTier1 (1)	Enterprise (2)	FlashTier0 (3)	FlashTier0
FlashTier2 (4)	FlashTier1	FlashTier1	FlashTier1
FlashTier0	FlashTier2	FlashTier2	FlashTier2
	Nearline (5)		Enterprise
	FlashTier0		Nearline

Notes:

1. The first tier of high capacity flash drives.
2. The tier that consists of drives with speeds of 10K RPM, 15K RPM, or a mixtures of 10K RPM and 15K RPM speeds.
3. The tier of high performance flash drives.
4. The second tier of high capacity flash drives.
5. The tier that consists of high-volume disks that are either SATA or SAS nearline drives.

Storage system configuration commands

Commands that are used to configure a storage system are provided for the DS command-line interface.

A storage system is a single, physical storage subsystem. A storage complex is a configuration of one or more storage systems that is managed by a management console. If you have one DS8000 machine, then you have a single storage complex containing a single storage system.

A storage image is a partitioning of a storage system that provides emulation of a virtual storage server. You can configure more than one storage image on a storage system.

The following storage system configuration commands are available:

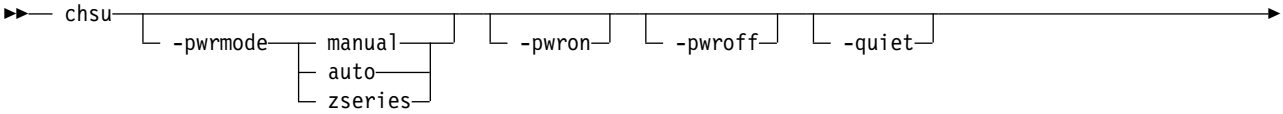
chsu Changes the description and name you have associated with a specified storage system. For DS8000, you can also change the remote power control mode on the storage system.

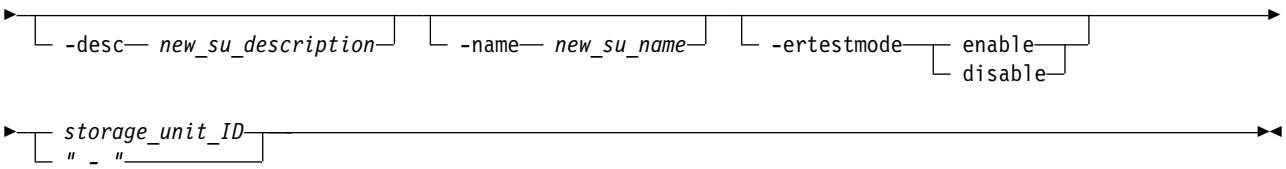
lssu Generates reports that allow you to view details about your storage systems.

showsu Generates reports that allow you to view details about your storage systems.

chsu

The **chsu** command modifies a storage system. You can also use this command to power on and power off a DS8000 storage system.





Parameters

-pwrmode manual | auto | zseries
(Optional) Sets a requested remote power control mode on the storage system.

manual
Indicates that the storage facility power-on and power-off sequences are done based on the manual power on and off controls.

auto
A storage facility power-on sequence is done when external power becomes available to the first rack of a storage facility. For example, a power-on sequence is done when standby power is activated to the remote power control cards.

zseries
Specifies that the power control mode is set to zSeries remote power control.

Note: Changing the power mode can take several minutes. Initiating a power-on or power-off request in manual mode can take up to 25 minutes. During a power-on or power-off request, access requests to the storage system might be queued. This queuing can result in a loss of response on other functions that access the storage system when accessed by the CLI.

-pwron
(Optional) Turns on power to the storage system. For DS8000, this parameter is valid if the control mode is set to manual and the switch is set to remote.

-pwroff
(Optional) Turns off power to the storage system. For DS8000, this parameter is valid if the control mode is set to manual and the switch is set to remote.

-quiet
(Optional) Turns off the modification confirmation prompt for this command.

-desc *new_su_description*
(Optional) Allows you to specify a description for the storage system. The description is limited to 256 byte or 128 double-byte characters.

-name *new_su_name*
(Optional) Allows you to specify a user-defined name for the storage system. This name is limited to 32 bytes or 16 double-byte characters.

-ertestmode enable | disable
(Optional) Enables or disables the Energy Report test mode.

enable
Energy Report readings are averaged over a 30-second period.

disable
Energy Report readings are averaged over a 5-minute period.

Note: IBM recommends disabling the Energy Report test mode at test completion.

storage_unit_ID | **-**
(Required) Accepts the fully qualified storage system ID. The storage system ID consists of manufacturer, machine type, and serial number. For example, IBM.2107-75FA120.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

Invoking the chsu command

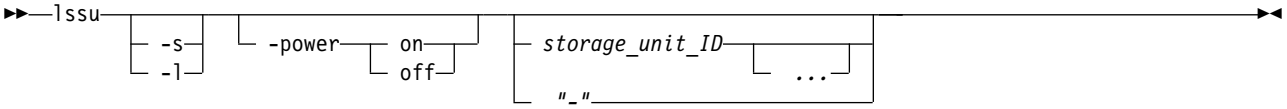
```
dscli> chsu -pwrmode manual IBM.2107-75FA120
```

The resulting output

```
Storage system IBM.2107-75FA120 successfully modified.
```

lssu

The **lssu** command displays a list of storage systems in a storage complex. You can use this command to review the status and other properties of each storage system in the list.



Parameters

-s (Optional) Displays only the storage system ID. You cannot use the **-l** and the **-s** parameters together.

-l (Optional) Displays default output plus the power mode and storage system description. You cannot use the **-l** and the **-s** parameters together.

-power on | off
(Optional) Displays only the storage systems in the specified power state.

storage_unit_ID ... | **-**
(Optional) Displays storage systems with the specified storage system IDs. A storage system ID includes manufacturer, machine type, and serial number. You must separate multiple IDs with a space between each ID.

Note: You cannot specify ID ranges.

The ellipsis (...) indicates that, optionally, you can specify multiple values. If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI list commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **lssu** command when you use the **-l** parameter.

Invoking the lssu command

```
dscli> lssu -l
```

The resulting output

Name	ID	Model	WWNN	Pw State	Pw Mode	Desc
SU 1	IBM.2107-75FA110	961	3007ACF 3012399 E0	On	Local	Test
SU 2	IBM.2107-75FA120	961	3007ACF 3045699 E0	On	Local	Produc-tion
SU 3	IBM.2107-75FA130	961	3007ACF 3078999 E0	On	Local	Backup

Report field definitions

Name
Indicates the user-defined name for each storage system that is found within the storage complex. This value is "-" if you did not assign a name to a storage system.

ID Indicates the storage system ID that consists of the manufacture name, machine type, and serial number. When the **-s** parameter is used, this is the only information that is displayed for the **lssu** command.

Model
Indicates the model number of the storage system.

WWNN
Indicates the World Wide Node Name for the listed storage system. This value is " -" if the WWNN is not known.

Pw State
Indicates the power status of the listed storage system. One of the following values is displayed:

On Indicates that the storage system power is on.

Off
Indicates that the storage system power is off.

Turning On
Indicates that the storage system power is turning on.

Turning Off
Indicates that the storage system power is turning off.

Power Exception
Indicates that storage system power is on, but online operation is not possible due to a power fault in one of the storage system frames.

Pw Mode
Indicates the power control mode in effect for the listed storage system. One of the following values is displayed:

Local
Indicates that the SMC local/remote switch is set to the local power control mode.

Remote SMC Manual
Indicates that the SMC local/remote switch is set to remote and that the power control mode is set to manual power control.

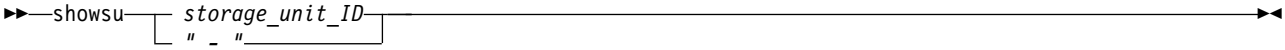
Remote SMC Auto
Indicates that the SMC local/remote switch is set to remote and that the power control mode is set to auto-power control.

Remote zSeries Power Control
Indicates that the SMC local/remote switch is set to remote and that the power control mode is set to zSeries remote power control.

Desc
Indicates the description that you assigned the storage system. This value is displayed as a " -" if no description is assigned.

showsu

The **showsu** command displays detailed properties of an individual storage system.



Parameters

storage_unit_ID
(Required) Specifies the storage system ID. A storage system ID consists of manufacturer, machine type, and serial number.

If you use the dash (-), the specified value is read from standard input. You cannot use the dash (-) while you are in the DS CLI interactive command mode.

Example

For this command and all other DS CLI show commands, the results are shown in table format to provide clarity. The actual reports do not display as tables.

The following table represents the headers that are displayed on the output report that is associated with the **showsu** command.

Invoking the showsu command

```
dscli> showsu IBM.2107-75FA120
```

The resulting output

Name	Desc	ID	Model	WWNN	Config
My Storage Unit	This is my DS Storage Unit	IBM.2107-75FA120	921	3007ACF 3012399 E0	4-way, single SFI

Pw State	Pw Mode	Reqpm	Processor Memory	MTS
On	Local	Remote SMC manual	1 GB	IBM.2424-75FA120

ER Test Mode	ER Recorded	ER Power Usage	ER Inlet Temp	ER I/O Usage	ER Data Usage
Disabled	2013-04-08T13:19:10-0700	5161	19.23	26	42

Report field definitions

Name
Indicates the name that you assigned for the designated storage system. This value is " - " if you did not assign a name to a storage system.

Desc
Indicates the description that you assigned for the designated storage system. This value is displayed as a " - " if no description is assigned.

ID Indicates the storage system ID that consists of the manufacture name, machine type, and serial number.

Model
Identifies the model number of the designated storage system.

WWNN
Indicates the World Wide Node Name for the listed storage system. This value is " - " if the WWNN is not known.

Config
Indicates the internal I/O interface configuration for the storage system. One of the following values is displayed:

Undefined
Indicates that the configuration is undefined. The value indicates that either this information is not supported for this storage system model, or a configuration upgrade is in progress that can cause the configuration option to change.

2-way, single SFI
Indicates that there are two processors and a single storage facility image.

4-way, single SFI
Indicates that there are four processors and a single storage facility image.

4-way, dual SFI
Indicates that there are four processors and dual storage facility images.

8-way, single SFI
Indicates that there are eight processors and a single storage facility image.

8-way, dual SFI
Indicates that there are eight processors and dual storage facility images.

Pw State
Indicates the power status of the listed storage system. One of the following values is displayed:

On Indicates that the storage system power is on.

Off
Indicates that the storage system power is off.

Turning On
Indicates that the storage system power is turning on.

Turning Off
Indicates that the storage system power is turning off.

Power® Exception
Indicates that storage system power is on, but online operation is not possible due to a power fault in one of the storage system frames.

Pw Mode
Indicates the power control mode in effect for the listed storage system. One of the following values is displayed:

Local
Indicates that the SMC local/remote switch is set to the local power control mode.

Remote SMC Manual
Indicates that the SMC local/remote switch is set to remote and that the power control mode is set to manual power control.

Remote SMC Auto
Indicates that the SMC local/remote switch is set to remote and that the power control mode is set to auto-power control.

Remote zSeries Power Control
Indicates that the SMC local/remote switch is set to remote and that the power control mode is set to zSeries remote power control.

Reqpm
Indicates the power control mode to apply when the local/remote switch is set to remote power control mode. One of the following values is displayed:

- Remote SMC Manual
- Remote SMC Auto
- Remote zSeries Power Control

Note: The default value is remote SMC Manual mode.

Processor Memory
Indicates the amount in decimal gigabytes (GB) of processor memory that is configured on the storage system.

MTS
(DS8000 models) Indicates the order type of the storage system. The order type varies according to the warranty periods associated with the storage system.

ER Test Mode
Indicates the state of the Energy Report test mode.

Disabled
The Energy Report test mode is disabled.

Enabled
The Energy Report test mode is enabled.

"-"
The dash (-) indicates that the Energy Report test mode is not supported on the system.

ER Recorded
The time stamp of the last successful Energy Report readings, in ISO8601 format. The format is *yyyy-MM-dd'T'HH:mm:ssZ*, where

- *yyyy* = the year
- *MM* = the month (01-12)
- *dd* = the day (01-31)
- *HH* = the hour (00-23)
- *mm* = the minutes (00-59)
- *ss* = the seconds (00-59)
- *Z* = the time zone offset from UTC [*-HHmm* | *+HHmm*]

Example: 2013-04-08T13:19:10-0700.

Unknown
Indicates that the Energy Report recording is supported on the system, but the value is not readable.

"-"
The dash (-) indicates that the Energy Report recording is not supported on the system.

ER Power Usage
The average total power usage, which is measured over a 5-minute period, in watts.

Unknown
Indicates that the Energy Report recording is supported on the system, but the value is not readable.

"-"
The dash (-) indicates that the Energy Report recording is not supported on the system.

ER Inlet Temp
The average inlet temperature, which is measured over a 5-minute period, in degrees Celsius.

Unknown
Indicates that the Energy Report recording is supported on the system, but the value is not readable.

"-"
The dash (-) indicates that the Energy Report recording is not supported on the system.

ER I/O Usage
The average I/O usage, which is measured over a 5-minute period, in KIOPS (1000 I/Os/second).

Unknown
Indicates that the Energy Report recording is supported on the system, but the value is not readable.

"-"
The dash (-) indicates that the Energy Report recording is not supported on the system.

ER Data Usage
The average data usage, which is measured over a 5-minute period, in MB/s (1,000,000 bytes/second).

Unknown
Indicates that the Energy Report recording is supported on the system, but the value is not readable.

"-"
The dash (-) indicates that the Energy Report recording is not supported on the system.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

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

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

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

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

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

© Copyright IBM Corp. 2004, 2019

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-IBMproducts. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

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

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

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/copytrade.shtml).

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

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

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

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

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

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

Index

A

account setup 17
AIX
  CLI installation considerations 1
  LIBPATH handling 1
AIX LIBPATH 6
alias volumes 436
applykey 38
audit report, console 245

C

chaccess 351
chauthpol 280
chckdvol 420
chcluster 46
checking command usage 28
chextpool 393
chfbvol 461
chhost 203
chhostconnect 200
chhostport 226
chkeymgr 317
chlcu 410
chlss 455
chpass 281
chpprc 127
chrank 380
chresgrp 271
chsession 153
chsestg 513
  space-efficient storage 513
chsi 533
chsp 359
chsu 545
chsyslogserver 265
chuser 282
chvolgrp 500
chzhyperlink 234
Cisco MDS 9216 Multilayer Fabric Switch 32, 122
CKD
  volume removal 440
CLI
  postinstallation 9
CLI commands on a command line 25
CLI commands, overview 35
CLI default configuration profile file 10
CLI help 29
CLI history function 28
CLI interactive mode 28
command line presentation 25
commands 354
  applykey 38
  chaccess 351
  chauthpol 280
  chckdvol 420
  chcluster 46
  chextpool 393
  chfbvol 461
commands *(continued)*
  chhost 203
  chhostconnect 200
  chhostport 226
  chkeymgr 317
  chlcu 410
  chlss 455
  chpass 281
  chpprc 127
  chprc 127
  chrank 380
  chresgrp 271
  chsession 153
  chsestg 513
  chsi 533
  chsp 359
  chsu 545
  chsyslogserver 265
  chuser 282
  chvolgrp 500
  chzhyperlink 234
  commitflash 54
  commitremoteflash 77
  cpauthpol 285
  diagsi 536
  dscli 161
  echo 162
  failbackpprc 109
  failoverpprc 112
  flagless parameters 35
  freezepprc 128
  helpmsg 164
  initckdvol 422
  initfbvol 464
  installsoftware 173
  lsaccess 352
  lsaddressgrp 409
  lsarray 371
  lsarraysite 365
  lsauthpol 285
  lsavailpprcport 96
  lsckdvol 424
  lscloudserver 41
  lscluster 48
  lsda 252
  lsddm 255
  lsextpool 394
  lsfbvol 465
  lsflash 59
  lsframe 260
  lsgmir 149
  lshba 261
  lshost 206
  lshostconnect 207
  lshostport 227
  lshosttype 212
  lshostvol 213
  lsioport 175
  lskey 39
  lskeygrp 317
  lskeymgr 322
commands *(continued)*
  lslcu 412
  lslss 457
  lsnetworkport 240
  lsportprof 183
  lspprc 115
  lspprcpath 98
  lsrank 381
  lsremoteflash 82
  lsresgrp 272
  lsserver 537
  lssession 156
  lssestg 516
  lssi 539
  lssoftware 174
  lsstgencl 263
  lssu 547
  lssyslogserver 267
  lsuser 286
  lsusergroup 288
  lsvolgrp 502
  lsvolinit 511
  lsvpn 362
  lswuiuser 289
  lszhyperlink 235
  lszhyperlinkport 236
  manageaccess 354
  manageckdvol 431
  managecloudserver 42
  manageextpool 397
  managefbvol 473
  managehostconnect 215
  managekeygrp 326
  managekeymgr 330
  managepwfile 290
  manageresckey 332
  manageresgrp 274
  mkaliasvol 435
  mkarray 375
  mkauthpol 292
  mkckdvol 436
  mkcloudserver 43
  mkcluster 50
  mkesconpprcpath 102
  mkextpool 398
  mkfbvol 478
  mkflash 64
  mkgmir 138
  mkhost 216
  mkhostconnect 217
  mkhostport 231
  mkkeygrp 334
  mkkeymgr 336
  mklcu 414
  mkpprc 122
  mkpprcpath 104
  mkrank 385
  mkreckey 338
  mkremoteflash 87
  mkresgrp 277
  mksession 155

© Copyright IBM Corp. 2004, 2019

commands *(continued)*
mksestg 522
mksyslogserver 268
mkuser 293
mkvolgrp 504
offloadauditlog 245
overview structure 35
parameters 35
pausegmir 139
pausepprc 129
resumegmir 140
resumepprc 131
resyncflash 55
resyncremoteflash 78
reverseflash 67
revertflash 71
revertremoteflash 90
rmarray 376
rmauthpol 295
rmckdvol 440
rmcloudserver 46
rmcluster 50
rmextpool 400
rmfbvol 482
rmflash 72
rmgmir 142
rmhost 220
rmhostconnect 220
rmhostport 231
rmkeygrp 338
rmkeymgr 339
rmlcu 416
rmpprc 133
rmpprcpath 106
rmrank 387
rmreckey 339
rmremoteflash 92
rmresgrp 278
rmsession 160
rmsestg 525
rmsyslogserver 270
rmuser 295
rmvolgrp 506
setauthpol 296
setenv 165
setflashrevertible 75
setioport 184
setnetworkport 239
setoutput 167
setremoteflashrevertible 94
setvpn 361
showaccess 357
showarray 376
showarraysite 369
showauthpol 305
showckdvol 441
showcluster 51
showenv 170
showextpool 401
showfbvol 484
showgmir 143
showgmircg 151
showgmiroos 152
showhost 221
showhostconnect 224
showhostport 232
showioport 186
commands *(continued)*
showkeygrp 340
showkeymgr 348
showlcu 417
showlss 458
shownetworkport 243
showpass 310
showrank 388
showresgrp 278
showsestg 527
showsi 541
showsp 363
showsu 549
showuser 311
showvolgrp 507
showzhyperlinkport 237
testauthpol 312
unfreezeflash 74
unfreezepprc 136
variables 35
ver 171
who 313
whoami 315
comments, sending ix
commitflash 54
commitremoteflash 77
configuration profile file 10
connection ID, caution 200
console audit report 245
controlling CLI page help 29
conventions
syntax v
cpauthpol 285
critical heavy mode 122
critical volume mode 122

D

default CLI settings 10
default directories 25
default report format 167
delim report format 167
diagsi 536
disableautoresync
failbackpprc 109
failoverpprc 112
mkpprc 122
resumepprc 131
DS CLI
default directories 25
first time login 25
interactive mode 28
limitations 32
OpenVMS messages 33
script mode 27
session limitation 32
single-shot mode 26
upgrade preparation 7
DS CLI command format 25
DS CLI exit codes 30
DS CLI installation
AIX considerations 1
default directories 1
HP Tru64 considerations 1
key aspects 1
System i considerations 1
DS CLI uninstall
Linux systems 8
dscli command 161

E

eam (extent allocation method)
lsckdvol 424
lsfbvol 465
mkckdvol 436
mkfbvol 478
showckdvol 441
showfbvol 484
EAV (extended addressing volume)
mkckdvol 436
echo command 162
encryption 330
encryption key 330
enhancements xi
exit codes, obtaining 30
extsize 385

F

failbackpprc 109
disableautoresync 109
tgtse 109
failoverpprc 112
disableautoresync 112
tgtse 112
feedback, sending ix
first time login 25
force parameter 109, 112
freezepprc 128

G

groups, users 288

H

help, CLI 29
helpmsg 164
HMC 357
host port commands 226
howaccess 357
HP-UX volume restriction 200

I

i5/OS
command format 26
i5/OS password file directory 290
IBM I
serial number reporting 541
incrementalresync 122
initckdvol 422
initfbvol 464
installation methods, overview of 3
installing the CLI 9
interactive mode
history function 28
setenv command 28

J

Java Virtual Machine Not Found Error 6

K

key group 330
Key Management Interoperability Protocol 322, 326, 336, 348
KMIP 322, 326, 336, 348

L

Linux
  lshostvol command considerations 213
Linux password file directory 290
Linux, uninstall DS CLI 8
lock account 282
lsaccess 352
lsaddressgrp 409
lsarray 371
lsarraysite 365
lsauthpol 285
lsavailpprcport 96
lsckdvol
  eam (extent allocation method) 424
  sam (standard allocation method) 424
lscloudserver 41
lscluster 48
lsda 252
lsddm 255
lsextpool 394
lsfbvol 465
  eam (extent allocation method) 465
  tse 465
lsflash 59
  tgtse, space-efficient volume 59
lsframe 260
lsgmir 149
lshba 261
lshost 206
lshostconnect 207
lshostport 227
lshosttype 212
lshostvol
  Linux considerations 213
  OpenVMS considerations 213
lsioport 175
lskey 39
lskeygrp 317
lskeymgr 322
lslcu 412
lslss 457
lsnetworkport 240
lsportprof 183
lspprc 115
lspprcpath 98
lsrank
  extsize 381
lsremoteflash 82
  tgtse, space-efficient volume 82
lsresgrp 272
lsserver 537
lssession 156
lssestg 516
lssestg *(continued)*
  space-efficient storage 516
lssi 539
lssoftware 174
lsstgencl 263
lssu 547
lssyslogserver 267
lsuser 286
lsusergroup 288
lsvolgrp 502
lsvolinit 511
lsvpn 362
lswuiuser 289
lszhyperlink 235
lszhyperlinkport 236
LUN sizes 478

M

manageaccess 354
manageckdvol 431
managecloudserver 42
manageextpool 397
managefbvol 473
managehostconnect 215
managekeygrp 326
managekeymgr 330
managepwfile 290
managereckey 332
manageresgrp 274
message types 33
mkaliasvol 435
mkarray 375
mkauthpol 292
mkckdvol 436
  eam (extent allocation method) 436
  EAV (extended addressing volume) 436
  tse 436
mkcloudserver 43
mkcluster 50
mkesconpprcpath 102
mkextpool 398
mkfbvol 478
  eam (extent allocation method) 478
  space-efficient storage 478
mkflash 64
  tgtse, space-efficient volume 64
mkgmir 138
mkhost 216
mkhostconnect 217
mkhostport 231
mkkeygrp 334
mkkeymgr 336
mklcu 414
mkpprc 122
  disableautoresync 122
  tgtse 122
mkpprcpath 104
mkrank 385
mkreckey 338
mkremoteflash 87
  tgtse, space-efficient volume 87
mkresgrp 277
mksession 155
mksestg 522
  space-efficient storage 522
mksyslogserver 268
mkuser 293
mkvolgrp 504

N

new features xi

O

offloadauditlog 245
OpenVMS
  DS CLI message considerations 33
  lshostvol command considerations 213
operational limitations 32

P

password file directories
  i5/OS 290
  LINUX 290
  UNIX 290
  Windows 290
pausegmir 139
pausepprc 129
postinstallation
  CLI 9
postinstallation tasks 9
processing tips
  3 site recovery 122
  cascade 122
  incrementalresync 122
  multiple volumes 122
  volume grouping method 122
publications
  ordering ix
  product v
  related v

R

RAID 6
  lsarray 371
  lsrank 381
  showarray 376
report formats
  controlling 28
  default 167
  delim 167
  setoutput command 167
  stanza 167
  XML 167
resumegmir 140
resumepprc 131
  disableautoresync 131
  tgtse 131
resyncflash 55
  tgtse, space-efficient volume 55
resyncremoteflash 78
  tgtse, space-efficient volume 78
reverseflash 67
  tgtse, space-efficient volume 67
revertflash 71
revertremoteflash 90

rmarray 376
rmauthpol 295
rmckdvol 440
rmcloudserver 46
rmcluster 50
rmextpool 400
rmfbvol 482
rmflash 72
  tgtse, space-efficient volume 72
rmgmir 142
rmhost 220
rmhostconnect 220
rmhostport 231
rmkeygrp 338
rmkeymgr 339
rmlcu 416
rmpprc 133
rmpprcpath 106
rmrank 387
rmreckey 339
rmremoteflash 92
  tgtse, space-efficient volume 92
rmresgrp 278
rmsession 160
rmsestg 525
  space-efficient storage 525
rmsyslogserver 270
rmuser 295
rmvolgrp 506

S

safeguarded copy
  lsckdvol 424
Safeguarded Copy
  managefbvol 473
safeguarded copy information,
  lsfbvol 465
Safeguarded Copy)
  showfbvol 484
script mode 27
serial number, obtaining 29
service action report, console 245
service audit log 245
setauthpol 296
setenv 165
setflashrevertible 75
  tgtse, space-efficient volume 75
setioport 184
setnetworkport 239
setoutput 167
setremoteflashrevertible 94
  tgtse, space-efficient volume 94
setup, account 17
setvpn 361
showaccess 357
showarray 376
showarraysite 369
showauthpol 305
showckdvol 441
  eam (extent allocation method 441
  tse 441
showcluster 51
showenv 170
showextpool 401
showfbvol 484
  eam (extent allocation method) 484
showfbvol *(continued)*
  tse 484
showgmir 143
showgmircg 151
showgmiroos 152
showhost 221
showhostconnect 224
showhostport 232
showioport 186
showkeygrp 340
showkeymgr 348
showlcu 417
showlss 458
shownetworkport 243
showpass 310
showrank 388
showresgrp 278
showsestg 527
  space-efficient storage 527
showsi 541
showsp 363
showsu 549
showuser 311
showvolgrp 507
showzhyperlinkport 237
single-shot mode
  command format, i5/OS 26
  command format, Windows 26
space-efficient storage
  chsestg 513
  lsfbvol 465
  lssestg 516
  mkfbvol 478
  mksestg 522
  rmsestg 525
  showsestg 527
space-efficient volume
  lsflash 59
  lsremoteflash 82
  mkflash 64
  mkremoteflash 87
  resyncflash 55
  resyncremoteflash 78
  reverseflash 67
  rmflash 72
  rmremoteflash 92
  setflashrevertible 75
  setremoteflashrevertible 94
stanza report format 167
status information, fixed block
  volume 465
stop the process
  Global Mirror 142
storage image ID, obtaining 29
syntax conventions v
System i
  CLI installation considerations 1
  interactive mode 28
  lunmapping 507
  postinstallation tasks 9
  script mode 27
  showvolgrp 507
  single-shot mode 26

T

testauthpol 312
tgtse
  failbackpprc 109
  failoverpprc 112
  lsflash 59
  lsremoteflash 82
  mkflash 64
  mkpprc 122
  mkremoteflash 87
  resumepprc 131
  resyncflash 55
  resyncremoteflash 78
  reverseflash 67
  rmflash 72
  rmremoteflash 92
  setflashrevertible 75
  setremoteflashrevertible 94
thaw operation 136
timeout connection 10
trademarks 554
tse
  lsckdvol 424
  lsfbvol 465
  showckdvol 441
  showfbvol 484
  space-efficient storage
    lsckdvol 424
tse, mksestg
  mkfbvol 478

U

unfreezeflash 74
unfreezepprc 136
uninstallation methods, overview of 8
uninstalling the DS CLI 8
UNIX password file directory 290
unlock account 282
user groups 288
user security 354
users 289

V

ver 171
volume format processing options 122
volume removal, CKD 440
volumes, alias 436

W

websites v
who 313
  engineering and service roles 313
whoami 315
Windows password file directory 290
write acceleration restriction 32, 122, 138

X

XML report format 167



Printed in USA

SC27-9562-00

