IBM DS8000 Series Version 7 Release 5

Command-Line Interface User's Guide



GC27-4212-08

Note

Before using this information and the product it supports, read the information in the Notices section.

This edition applies to version 7, release 5, Modification 2 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, 2016.

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

Contents

About this book	
Who should use this book	
Conventions and terminology	
Syntax diagram conventions	
Publications and related information	
Ordering IBM publications	
Summary of Changes	
Chapter 1. Introduction to the DS8000	
series	
Overview of the DS8000 series	
Chapter 2. Installing, upgrading, and	
uninstalling the DS CLI	
Operating systems that support the DS CLI 3	
Installing the DS CLI	
DS CLI preinstallation information 4	
Mounting the DS CLI installation CD 6	
Installing the DS CLI using the graphical mode 7	
Installing the DS CLI using the console mode 10 Installing the DS CLI using the unattended	
(silent) mode	
Error	
Upgrading the DS CLI on your system 16	
Uninstalling the DS CLI	
Uninstalling the DS CLI from your system using	
graphical mode	
Uninstalling the DS CLI using unattended (silent)	
mode	
Uninstalling the DS CLI using the console mode 18	
Uninstalling the DS CLI from a System i model 18 Running and configuring the DS CLI 20	
Creating a default CLI profile	
Setting up user accounts using the DS CLI	
Activating your machine and feature licenses	
using the DS CLI	
OpenVMS system integration	
Chapter 3. Running the DS CLI 35	
Logging into the DS CLI	
Using the DS CLI single-shot command mode	
Using the DS CLI script command mode 37	
Using the DS CLI interactive command mode	
(history and reports)	
Obtaining the serial (storage image ID) number	
using the DS CLI	
DS CLI command help	
Obtaining and interpreting DS CLI exit codes 40 DS CLI operational limitations	
DS CLI operational limitations	
server	

Chapter 4. CLI commands				. 45
About CLI commands				. 45
Understanding the syntax diagrams .				
Common command flags				
Framework commands				. 48
dscli				
dscli				. 50
exit				
help				
helpmsg				
quit	•	• •	·	. 52
quit	•	• •	•	
setenv	•	• •	•	. 55
showenv	•	• •	·	. 55
ver	•		•	. 59
Security commands.	•	• •	·	. 61
User account and security commands Data encryption and security comman User access security commands			•	. 61
Data encryption and security comman	nds		•	. 94
User access security commands	•		•	. 116
Internet Protocol Security (IPSec) com				
Application key commands				. 137
applykey				. 137
lskey				. 138
Physical resource information command	s.			. 139
Ísda				. 140
lsframe				. 142
lsstgencl				
lshba				
leddm	·	•	·	148
lsddm	اد	·	•	152
chsp	15.	•	·	152
setvpn	·	•	·	152
lsvpn	·	·	·	. 154
showsp	·	•	·	. 155
Storage unit configuration commands .	·	·	·	. 157
chsu	·	•	·	. 158
lssu	•	•	·	. 159
showsu	·	•	·	. 161
Storage image configuration commands		•	•	. 164
chsi		•	•	. 165
diagsi			•	. 167
lsserver				. 168
lssi				. 170
showsi				. 172
I/O port and host connection configurat				
commands				. 176
Storage image I/O port commands .				. 176
Host connection commands				. 192
		:		. 212
				. 212
Array site specific commands				· 212 · 218
Array specific commands				
Rank specific commands				. 226
Extent pool specific commands				. 239
Address group specific commands .	. ·	•		. 255
Logical control unit specific command				. 256
CKD logical volume specific comman	ds	•	•	. 267

Logical subsystem specific commands	300
Fixed block logical volume specific commands	306
Volume group specific commands	342
Advanced operation commands	352
Space-efficient storage commands	355
I/O Priority Management commands	376
lsperfgrp	376
lsperfgrprpt	378
lsperfrescrpt	380
Resource Group commands	382
chresgrp	382
lsresgrp	383
manageresgrp	384
mkresgrp	388
rmresgrp	389
showresgrp	389
Copy Services commands	391
FlashCopy commands	391
Remote FlashCopy commands	415
Remote Mirror and Copy path commands.	435
Remote Mirror and Copy commands	448
Global Mirror commands	478
Global Mirror session commands	495
Offload file commands	503
offloadauditlog	504
offloadfile	506

Chapter 5. Configuring access and

security	. 511
Internet Protocol Security (IPSec) tasks	. 511
Creating an IPSec connection configuration file	e 511
Creating a PSK IPSec connection	. 512
Creating a public key IPSec connection	. 513
Enabling or disabling an IPSec connection.	. 514
Listing IPSec connections	. 514
Export an IPSec connection file	. 515
Starting or stopping the IPSec server	. 516
Importing certificates	. 516
Listing certificates	. 517

Chapter 6. Configuring and managing

logical storage 51	3
Configuring new fixed block storage using the DS	
CLI	9
Creating extent pools for fixed block volumes	
using the DS CLI	9
Creating arrays for fixed block volumes using	
the DS CLI \ldots \ldots \ldots \ldots \ldots \ldots 52	21
Creating a rank using the DS CLI	22
Creating fixed block volumes using the DS CLI 52	
Creating LUN volumes for System i models 52	
Correcting a fixed block configuration error 52	26
Creating fixed block volume groups using the	
DS CLI	
Creating a volume group for System i models 52	28
Configuring Fibre Channel I/O ports using the	
DS CLI	29
Creating SCSI host port connections using DS	
CLI	30

Configuring new count key data storage using DS	
CLI.	532
Creating count key data extent pools using the	00-
DS CLI	532
Creating arrays for CKD volumes using the DS	
CLI	534
Creating a rank for CKD volumes using the DS	
CLI	535
Creating logical control units for CKD volumes	
using DS CLI	536
Creating count key data volumes using the DS	
	537
Correcting a CKD volume configuration error	538
Configuring Fibre Channel I/O ports using the DS CLI	539
	540
	540
0	542
Modifying an extent pool	542
Viewing extent pool properties and performance	012
metrics	543
Deleting extent pools from a storage	
configuration	543
0	544
	545
Viewing array status .	545
Removing arrays from a storage configuration	
	546
	547
Modifying a rank	548
	549
	549
Correcting a rank-related configuration error	549
Removing ranks from a storage configuration	550
5 0 0	551
Viewing logical control unit status	552
Viewing properties for one logical control unit	552
Removing logical control units from a CKD	553
storage configuration	555
Chapter 7. Copy Services functions	555
FlashCopy functions	555
	555
Creating a persistent FlashCopy relationship	556
Viewing information about FlashCopy	000
	556
	557
	558
	560
Reversing a FlashCopy relationship	560

Creating a FlashCopy target volume on an

Discarding changes to FlashCopy target

existing Metro Mirror source volume 563

Metro Mirror functions	. 50	66
Displaying the status of established paths	. 50	
Displaying the WWNN of a storage unit	. 50	
Creating remote mirror and copy paths	. 50	67
Correcting a path-related configuration error		68
Removing paths	. 50	
Creating a Metro Mirror relationship	. 50	
Creating a Metro Mirror consistency group .	. 52	
Resuming a Metro Mirror relationship	. 52	
Pausing a Metro Mirror relationship	. 52	
Creating a Global Copy relationship	. 52	
Deleting a Metro Mirror relationship	. 52	
Modifying logical subsystem timeout values .	. 52	74
Defining a path that has the consistency option		
enabled	. 52	
Monitoring Remote Mirror and Copy paths .	. 52	
Running a recovery failback operation	. 52	
Running a recovery failover operation	. 52	78
Viewing information about Metro Mirror		
relationships	. 52	78
Converting Global Copy volume pairs to		
synchronous	. 52	79
Determining which I/O ports are available for		
paths	. 52	79
Deleting a single volume Metro Mirror		
relationship	. 58	80
Copy Services functions across a 2105 and		
DS8000 or DS6000 model	. 58	81
Creating a Metro Mirror volume pair between a		
DS8000 model or a DS6000 model and a 2105	52	82
	. 0	2
Global Mirror functions		84
Adding volumes to a session (Global Mirror)	. 58	
	. 58	84
Adding volumes to a session (Global Mirror)	. 58	84 84
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global	. 58 58	84 84
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session	. 58 58	84 84 85
Adding volumes to a session (Global Mirror)Modifying the tuning parameters of a GlobalMirror sessionModifying the topology of a Global MirrorsessionViewing a Global Mirror session	. 58 58 . 58	84 84 85 86
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session	. 58 58 . 58	84 84 85 86 88
Adding volumes to a session (Global Mirror)Modifying the tuning parameters of a GlobalMirror sessionModifying the topology of a Global MirrorsessionSessionViewing a Global Mirror sessionQuerying Global Mirror processingPausing Global Mirror processing	. 58 . 58 . 58 . 58	84 84 85 86 88 88
Adding volumes to a session (Global Mirror)Modifying the tuning parameters of a GlobalMirror sessionModifying the topology of a Global MirrorsessionSessionViewing a Global Mirror sessionQuerying Global Mirror processing	. 58 . 58 . 58 . 58 . 58	84 84 85 86 88 88 88
Adding volumes to a session (Global Mirror)Modifying the tuning parameters of a GlobalMirror sessionModifying the topology of a Global MirrorsessionsessionViewing a Global Mirror sessionQuerying Global Mirror processingPausing Global Mirror processingResuming Global Mirror processing	. 58 . 58 . 58 . 58 . 58 . 58 . 58 . 58	84 84 85 86 88 88 88 89 89
Adding volumes to a session (Global Mirror)Modifying the tuning parameters of a GlobalMirror sessionModifying the topology of a Global MirrorsessionsessionViewing a Global Mirror sessionQuerying Global Mirror processingPausing Global Mirror processingResuming Global Mirror processingStarting Global Mirror processing	 . 58 	84 84 85 86 88 88 88 89 89
Adding volumes to a session (Global Mirror)Modifying the tuning parameters of a GlobalMirror sessionModifying the topology of a Global MirrorsessionSessionViewing a Global Mirror sessionQuerying Global Mirror processingPausing Global Mirror processingResuming Global Mirror processingStarting Slobal Mirror Processing <td> . 58 . 59 . 59 . 59 . 59 . 59 . 59 </td> <td> 84 84 85 86 88 89 89 90 90 90 </td>	 . 58 . 59 . 59 . 59 . 59 . 59 . 59 	 84 84 85 86 88 89 89 90 90 90
Adding volumes to a session (Global Mirror)Modifying the tuning parameters of a GlobalMirror sessionModifying the topology of a Global MirrorsessionsessionViewing a Global Mirror sessionQuerying Global Mirror processingPausing Global Mirror processingResuming Global Mirror processingStarting Global Mirror processingStarting Global Mirror processingEnding Global Mirror processingMirror processingStarting Global Mirror processingStarting Global Mirror processingEnding Global Mirror processingStarting Slobal Mirror ProcessingStart	 . 58 . 59 . 59 . 59 . 59 . 59 . 59 	84 84 85 86 88 88 89 90 90 90
Adding volumes to a session (Global Mirror)Modifying the tuning parameters of a GlobalMirror sessionModifying the topology of a Global MirrorsessionsessionViewing a Global Mirror sessionQuerying Global Mirror processingPausing Global Mirror processingResuming Global Mirror processingStarting Global Mirror processingStarting Global Mirror processingEnding Global Mirror processingStarting Global Mirror processingSetting Up the Global Mirror Environment.	 . 58 . 59 . 59 . 59 . 59 . 59 . 59 	84 84 85 86 88 88 89 90 90 91 91
Adding volumes to a session (Global Mirror)Modifying the tuning parameters of a GlobalMirror sessionModifying the topology of a Global MirrorsessionsessionViewing a Global Mirror sessionQuerying Global Mirror processingPausing Global Mirror processingResuming Global Mirror processingStarting Global Mirror processingEnding Global Mirror processingStarting Global Mirror Processing	 . 5% . 5%<td>84 84 85 86 88 88 89 90 90 91 91</td>	84 84 85 86 88 88 89 90 90 91 91
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session	 . 5% . 5%<td>84 84 85 86 88 88 89 90 90 91 91</td>	84 84 85 86 88 88 89 90 90 91 91
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session	 58 58 58 58 58 58 59 59 59 59 59 59 59 	84 85 86 88 89 90 91 91 91 96
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session	 . 5% . 5%<td>84 85 86 88 89 90 91 91 91 96</td>	84 85 86 88 89 90 91 91 91 96
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session	 . 58 . 58 . 58 . 58 . 58 . 58 . 59 . 59 . 59 . 59 . 59 . 59 . 60 	84 84 85 86 88 89 90 91 91 91 91 91 91 91
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session	 58 58 58 58 58 58 59 59 59 59 59 59 59 	84 84 85 86 88 89 90 91 91 91 91 91 91 91
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session	 . 53 . 53 . 53 . 53 . 53 . 54 . 54<td>84 84 85 86 88 89 90 91 91 91 91 96)3 04</td>	84 84 85 86 88 89 90 91 91 91 91 96)3 04
Adding volumes to a session (Global Mirror)Modifying the tuning parameters of a GlobalMirror sessionModifying the topology of a Global MirrorsessionSessionViewing a Global Mirror processingQuerying Global Mirror processingPausing Global Mirror processingResuming Global Mirror processingStarting Global Mirror processingEnding Global Mirror processingStarting Global Mirror processingStarting Global Mirror processingEnding Global Mirror processing (script mode)Ending Global Mirror processing (no script)Setting up the Global Mirror EnvironmentRemoving a Global Mirror environmentEnding Global Mirror processing when a disasterusing Global Mirror processing when a disasteroccursChecking Global Mirror transaction status in adisaster situation	 . 58 . 58 . 58 . 58 . 58 . 58 . 59 . 60 	84 84 85 86 88 89 90 91 91 91 91 96)3 04
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session	 53 54 <	84 85 86 88 89 90 91 91 91 96 D3 04
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session Modifying the topology of a Global Mirror session Viewing a Global Mirror session Querying Global Mirror processing Pausing Global Mirror processing Resuming Global Mirror processing Starting Global Mirror processing Ending Global Mirror processing Ending Global Mirror processing Starting Global Mirror processing Ending Global Mirror processing (script mode) Ending Global Mirror processing (no script) Setting up the Global Mirror Environment. Removing a Global Mirror environment Removing a Global Mirror processing when a disaster	. 58 58 . 58 . 58 . 58 . 58 . 59 . 59 . 59 . 59 . 59 . 59 . 59 . 59	84 85 86 88 89 90 91 91 91 96 D3 04
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session	 53 53 53 53 53 53 54 <	84 84 85 86 88 89 90 91 91 91 96)3 04 04 05
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session	 53 54 <	84 84 85 86 88 89 90 91 91 91 96)3 04 04 05
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session Modifying the topology of a Global Mirror session Viewing a Global Mirror session Querying Global Mirror processing Pausing Global Mirror processing Pausing Global Mirror processing Starting Global Mirror processing Starting Global Mirror processing Ending Global Mirror processing Ending Global Mirror processing Starting Global Mirror processing Ending Global Mirror processing (no script) . Setting up the Global Mirror Environment. Removing a Global Mirror environment Removing a Global Mirror renvironment Chapter 8. Recovering from a disaster occurs Checking Global Mirror processing when a disaster occurs Initiating failover processing for B volumes to A volumes	 53 54 <	84 85 86 88 89 90 91 91 91 91 91 91 91 91 91 91 91 91 91
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session Modifying the topology of a Global Mirror session Viewing a Global Mirror session Querying Global Mirror processing Pausing Global Mirror processing Resuming Global Mirror processing Starting Global Mirror processing Ending Global Mirror processing Ending Global Mirror processing Starting Global Mirror processing Ending Global Mirror processing (no script) . Setting up the Global Mirror Environment. Removing a Global Mirror environment Removing a Global Mirror processing when a disaster occurs Chapter 8. Recovering from a disaster occurs Initiating Global Mirror transaction status in a disaster situation disaster situation Initiating failover processing for B volumes to A volumes <	 53 53 53 53 53 53 54 <	84 85 86 88 89 90 91 91 91 91 91 91 91 91 91 91 91 91 91
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session Modifying the topology of a Global Mirror session Viewing a Global Mirror session Querying Global Mirror processing Pausing Global Mirror processing Resuming Global Mirror processing Starting Global Mirror processing Ending Global Mirror processing Ending Global Mirror processing Ending Global Mirror processing Starting Global Mirror processing (no script) . Setting up the Global Mirror Environment. Removing a Global Mirror environment Removing a Global Mirror processing when a disaster occurs Chapter 8. Recovering from a disaster occurs Initiating Global Mirror transaction status in a disaster situation disaster situation Initiating failover processing for B volumes to A volumes <	 53 54 <	84 85 86 88 89 90 91 91 91 96 03 04 04 05 06 08
Adding volumes to a session (Global Mirror) Modifying the tuning parameters of a Global Mirror session Modifying the topology of a Global Mirror session Viewing a Global Mirror session Querying Global Mirror processing Pausing Global Mirror processing Resuming Global Mirror processing Starting Global Mirror processing Ending Global Mirror processing Ending Global Mirror processing Starting Global Mirror processing Ending Global Mirror processing (no script) . Setting up the Global Mirror Environment. Removing a Global Mirror environment Removing a Global Mirror processing when a disaster occurs Chapter 8. Recovering from a disaster occurs Initiating Global Mirror transaction status in a disaster situation disaster situation Initiating failover processing for B volumes to A volumes <	 53 54 <	84 85 86 88 89 90 91 91 91 96 03 04 04 05 06 08

Using fast reverse restore processing to create	
consistency	610
Waiting for the background copy to complete	
Reestablishing the FlashCopy relationships between	
· · ·	611
Preparing to reinstate production at the local site	612
Resynchronizing the volumes	
Querying, quiescing, and re-querying	
Reestablishing remote mirror and copy paths (site	
A to site B)	615
Running Global Copy failover processing to the A	0.20
volumes	617
Running Global Copy failback processing for the A	017
volumes	617
Resuming Global Mirror processing at site A	
incoming clocar mintor processing at site H	010

Chapter 9. Recovery scenarios for planned and unplanned outages using Metro/Global Mirror 621

Chapter 10. Command-line interface

scenarios	685
Modifying fixed block volume groups	685
Deleting data storage configurations	685
Deleting a fixed block data storage	
configuration	686
Deleting a count key data storage configuration	689
Processing remote FlashCopy (inband) transactions	691
Metro Mirror test scenario: failback operation from	
local to remote site	693
Allowed remote mirror and copy volume pair	
1, 1	694
Resource Groups	695
	697
Appendix. Archived CLI information 7	701

Notices				-	703
Trademarks					. 704
Homologation statement					. 704

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 DS8000 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 series.

Conventions and terminology

Different typefaces are used in this guide to show emphasis, and various notices are used to highlight key information.

Typeface	Description
Bold	Text in bold represents menu items.
bold monospace	Text in bold monospace represents command names.
Italics	Text in italics is used to emphasize a word. In command syntax, it is used for variables for which you supply actual values, such as a default directory or the name of a system.
Monospace	Text in monospace identifies the data or commands that you type, samples of command output, examples of program code or messages from the system, or names of command flags, parameters, arguments, and name-value pairs.

The following typefaces are used to show emphasis:

These notices are used to highlight key information:

Notice	Description
Note	These notices provide important tips, guidance, or advice.
Important	These notices provide information or advice that might help you avoid inconvenient or difficult situations.
Attention	These notices indicate possible damage to programs, devices, or data. An attention notice is placed before the instruction or situation in which damage can occur.

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 DS8000 series command syntax.

Table 1. Command syntax conventions

Syntax convention	Description	Example
	A command is the first word or set of consecutive characters.	help

Table 1. Command syntax conventions (continued)

Syntax convention	Description	Example
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 <i>italic</i> typeface and can use capitalized letters in the character string to aid in reading comprehension.	timeout_in_sec
Vertical bar (1)	Mutually exclusive options are separated by a vertical bar (1).	[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 (www.ibm.com/support/knowledgecenter/ ST8NCA/product_welcome/ds8000_kcwelcome.html) contains all of the 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 (www.ibm.com/shop/publications/order/). Publications without an order number are available on the documentation CD or can be downloaded here.

Title	Description	Order number
DS8870 Introduction and Planning Guide	the product and technical concepts for	V7.5 GC27-4209-11 V7.4 GC27-4209-10 V7.3 GC27-4209-09 V7.2 GC27-4209-08 V7.1 GC27-4209-05
	system.	V7.0 GC27-4209-02

Table 2. DS8000 series product publications

Table 2. DS8000 series	product publications	(continued)
------------------------	----------------------	-------------

Title	Description	Order number
DS8800 and DS8700 Introduction and Planning Guide	This publication provides an overview of the product and technical concepts for DS8800 and DS8700. It also describes ordering features and how to plan for an installation and initial configuration of the storage system.	V6.3 GC27-2297-09 V6.2 GC27-2297-07
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 System $z^{\textcircled{B}}$ hosts.	V7.5 GC27-4210-04 V7.4 GC27-4210-03 V7.2 GC27-4210-02 V7.1 GC27-4210-01 V7.0 GC27-4210-00 V6.3 GC27-2298-02
IBM Storage System Multipath Subsystem Device Driver User's Guide	This publication provides information regarding the installation and use of the Subsystem Device Driver (SDD), Subsystem Device Driver Path Control Module (SDDPCM), and Subsystem Device Driver Device Specific Module (SDDDSM) on open systems hosts.	Download
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.	V7.5 GC27-4212-06 V7.4 GC27-4212-04 V7.3 GC27-4212-03 V7.2 GC27-4212-02 V7.1 GC27-4212-01 V7.0 GC27-4212-00 V6.3 GC53-1127-07
Application Programming Interface Reference	This publication provides reference information for the DS8000 Open application programming interface (DS Open API) and instructions for installing the Common Information Model Agent, which implements the API.	V7.3 GC27-4211-03 V7.2 GC27-4211-02 V7.1 GC27-4211-01 V7.0 GC35-0516-10 V6.3 GC35-0516-10
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.	V1.0 SC27-8502-00

Table 3.	DS8000 series	warranty, notices	and licensina	publications
10010 0.	200000 001100	manany, nonooc	, and noononig	pablicationic

Title	Order number
Warranty Information for DS8000 series	See the DS8000 Publications CD
IBM Safety Notices	Search for G229-9054 on the IBM Publications Center website
IBM Systems Environmental Notices	http://ibm.co/ 1fBgWFI

Table 3. DS8000 series warranty, notices, and licensing publications (continued)

Title	Order number
International Agreement for Acquisition of Software Maintenance (Not all software will offer Software Maintenance under this agreement.)	http://ibm.co/ 1fBmKPz
License Agreement for Machine Code	http://ibm.co/ 1mNiW1U
Other Internal Licensed Code	http://ibm.co/ 1kvABXE
International Program License Agreement and International License Agreement for Non-Warranted Programs	Download

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 publications

Listed here are the IBM Redbooks[®] publications, technical papers, and other publications that relate to DS8000 series.

Table 4.	DS8000	series	related	publications

Title	Description
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 Tivoli [®] Storage Productivity Center online product documentation (www.ibm.com/support/ knowledgecenter/SSNE44/)	This online documentation provides information about Tivoli Storage Productivity Center, which you can use to centralize, automate, and simplify the management of complex and heterogeneous storage environments includingDS8000 storage systems and other components of your data storage infrastructure.

Related websites

View these websites to get more information about DS8000 series.

Table 5. DS8000 series related websites

Title	Description
IBM website (ibm.com [®])	Find more information about IBM products and services.
IBM Support Portal website (www.ibm.com/storage/support)	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 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.

Table 5. DS8000 series related websites (continued)

Title	Description
IBM Storage SAN (www.ibm.com/systems/ storage/san)	Find information about IBM SAN products and solutions, including SAN Fibre Channel switches.
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.
DS8700 Code Bundle Information website (www.ibm.com/support/ docview.wss?uid=ssg1S1003593)	Find information about code bundles for DS8700. See section 3 for web links to SSD information. The version of the currently active installed code bundle now displays with the DS CLI ver command when you specify the -1 parameter.
DS8800 Code Bundle Information website(www.ibm.com/support/ docview.wss?uid=ssg1S1003740)	Find information about code bundles for DS8800. See section 3 for web links to SSD information. The version of the currently active installed code bundle now displays with the DS CLI ver command when you specify the -1 parameter.
DS8870 Code Bundle Information website (www.ibm.com/support/ docview.wss?uid=ssg1S1004204)	Find information about code bundles for DS8870. See section 3 for web links to SSD information.The version of the currently active installed code bundle now displays with the DS CLI ver command when you specify the -1 parameter.

Ordering IBM publications

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

The IBM Publications Center website (www.ibm.com/shop/publications/order/) offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download at no charge. You can also order publications. The IBM Publications Center website displays prices in your local currency.

Sending comments

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

To submit any comments about this publication or any other DS8000 series documentation:

Send your comments by email to starpubs@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 7 Release 5, Modification 2 introduces the following new features. For DS8000 series information, see the IBM DS8000 Introduction and Planning Guide for your specific model.

Version 7, Release 5, Modification 2

This table provides the current technical changes to the DS8000 series (as of July 2016).

Function	Description
Modified commands	• Added the loginprompt and promptmessage parameters to the chsp command. These parameters were also added to the report definitions of the showsp command.
	• Added the spidfstate parameter to the following commands:
	– showlcu
	– showlss
	• Added the spidfdisable parameter to the following commands:
	– manageckdvol
	– managefbvol
	• Added the config parameter to the offloadfile command.
	 Modified the port parameter to serverport in the mkkeymgr and lskeymgr commands. The parameter was also modified in the report definitions of the lskeymgr command.

Version 7, Release 5, Modification 1

This table provides the current technical changes to the DS8000 series (as of June 2015).

Function	Description
mksestg command	Updated the mksestg command to reflect the following text under the repcap parameter:
	 The minimum repository capacity that can be created is as follows: gb = 16 GiB blocks = 33 554 432 blocks, which is equivalent to 16 GiB cyl = 17 808 cylinders
OpenVMS	The information is included in Appendix. Archived CLI information.

Function	Description
Modified commands	 setauthpol command - Added the enable, disable, and setlocaladmin parameters. showauthpol command - Added the localAdmin field to the report definition. lsioport command - Added the CmdRetries and TransferReady fields to the report definitions. Included read diagnostic parameters TxPower , RxPower, TransceiverTemp, SupplyVolt, TXBiasCurrent, ConnectorType, TxType, CurrentSpeed, FECStatus, UncorrectedBitErr, and CorrectedBitErr fields to the report definitions. showioport command - Added the parameters included in the lsioport command above to the report definitions.

This table provides the current technical changes to the DS8000 series (as of May 22, 2015).

Version 7, Release 4, Modification 1

This table provides the current technical changes to the DS8000 series (as of January 2015).

Function	Description
CLI default profile	Updated information, including the list of profile variables that you can use to create a DS CLI profile. For information, see "Creating a default CLI profile" on page 21.

Version 7, Release 4

This table provides the current technical changes and enhancements to the DS8000 series (as of December 5, 2014).

Function	Description
New commands	 manageextpool command - Starts a process to initiate a change on extent pools. chpprc command - Modifies the characteristics of an existing Remote Mirror and Copy relationship.

Function	Description
Modified commands	• Isrank and showrank commands - Added the marray parameter.
	• 1shostconnect and showhostconnect commands - Added the host field to the report definitions.
	 showextpool command - Added the etmigpauseremain and etmonpauseremain and etmonitorreset fields to the report definitions.
	 managefbvol and manageckdvolcommands - Added the etmonpause and etmonresume and etmonreset parameters.
	• showfbvol and showckdvol commands - Added the Assign Pending Hardware to the tierassignstatus parameter. Added etmonpauseremain and etmonitorreset fields to the report definitions.
	• mkflash, resyncflash, reverseflash, mkremoteflash, resyncremoteflash and lsflash commands - Added the multinc parameter.
	• failoverpprc and lspprc commands - Added the multtgt parameter.
	 mksession and chsession commands - Added the remotedev storage_image_ID variable and the volpair parameter.
	• Issession command - Added the SecondaryVolume field to the report definitions.
	• managekeymgr and managekeygrp commands - Added the testaccess parameter.
	• lskeymgr and showkeymgr commands - Added the critical and not_normal field to the report definitions.
	 showkeymgr and showkeygrp commands - Added the access parameter. Added the lastaccess, lastsuccess, lastfailure, grpstatus, and mgrstatus fields to the report definitions.
	• 1skeymgr command - Added the grpstatus and mgrstatus field to the report definitions.
	• chvolgrp command - Added the safe parameter.

Function	Description
IBM Knowledge Center	To view DS CLI commands reference and usage information in Knowledge Center, use the search or filtering functions, or find it in the navigation by clicking System Storage > Disk systems > Enterprise Storage Servers > DS8000 and see Reference in the navigation. Go to the IBM Knowledge Center website to learn more.
Java 6	Attention: The DS CLI now requires the Java 6 platform, or later, to be installed before the installation of the DS CLI.
Lightweight on-demand dump (ODD)	ODD provides support to capture data for analysis with no impact to host applications. For information, see the diagsi command

Function	Description
Modified commands	• lsstgencl command - Added the DA_Pair_Type field to the report definitions.
	• 1sda command - Added the WWN field to the report definitions.
	 1sddm command - Added the Flash option to the diskclass parameter and the Flash field to the report definitions.
	• lsarrysite command - Added the Flash option to the diskclass parameter and the Flash field to the report definitions.
	• showarraysite command - Added the Flash field to the report definitions.
	• Isarray command - Added the Flash field to the report definitions.
	• showarray command - Added the Flash field to the report definitions.
	• showckdvol command - Added the sysplex field to the report definitions.
	• showfbvol command - Added the sysplex field to the report definitions.
	• diagsi command - Added the odd and oddlite options to the action parameter.
	 ver command - Added the Bundle Version and HMC DSCLI fields to the report definitions.

Function	Description
Modified commands	• chsestg command - Added reptype parameter.
	• chsu command - Added ertestmode parameter.
	• dscli command - Added port parameter.
	• 1saccess command - Added vpn field.
	• lsarraysite command - Added diskclass and diskrpm parameters.
	• Issesstg command - Added reptype and opratio fields.
	 managekeygrp command - Added updatecert value to action parameter.
	• managepwfile command - Added mcall parameter.
	• mkkeymgr command - Modified description and added cert parameter.
	• mksestg command - Added reptype parameter.
	• rmsestg command - Added reptype parameter.
	• setenv command - Added banner_date parameter.
	• showenv command - Added bannerDate value.
	• showkeygrp command - Added certificate field.
	 showsestg command - Added reptype and opratio fields.
	• showsu command - Added Energy Report Test Mode, ER Recorded, ER Power [®] Usage, ER Inlet Temp, ER I/O Usage, and ER Data Usage fields.
	• who command - Added protocol parameter and protocol field.
New commands	• manageaccess command - Manages the minimum security protocol access settings of an HMC for all communications to and from the DS8000 system.
	• managekeymgr command - Allows you to manage an existing encryption key server.
	• showaccess command - Displays the access properties of a specified HMC.
	• showkeymgr command - Displays detailed properties of a specified key manager entry.

Function	Description
Removed command	 setrmpw - Changes the IBM Tivoli Storage Productivity Center Replication Manager password. Information about this command is now available in the <i>Archived CLI Information</i> section of the IBM System Storage DS8000 Information Center.
Default CLI profile creation	Added port variable to specify which port the DSCLI should use when connecting to the DS8000 system.

Function	Description
Modified commands	chsi command - Added ethmtmode and etccmode parameters.
	• showsi command - Added ETHMTMode and ETCCMode fields.
	• 1spprc command - Added firstpass and type parameters.
	• mkfbvol, chfbvol, and lsfbvol commands - Added types 050 and 099 for IBM i.
	• showvolgrp command - Added etccmap parameter.
	 showckdvol command - Added tier and pathgrp parameters. Added the following new fields: tierassignstatus, tierassignerror, tierassignorder, tierassigntarget, %tierassigned.
	• showfbvol command - Added tier , pathgrp , and reserve parameters. Added the following new fields: tierassignstatus, tierassignerror, tierassignorder, tierassigntarget, %tierassigned.
	• mkflash command - Added resetreserve parameter.
	• resyncflash command - Added resetreserve parameter.
	• resyncremoteflash command - Added resetreserve parameter.
	• reverseflash command - Added resetreserve parameter.
	• mkremoteflash command - Added resetreserve parameter.
	• showl cu command - Added sfstate parameter.
	• showlss command - Added sfstate parameter.
	 manageckdvol command - Added tier and assignorder parameters. Added sfdisable, tierassign, and tierunassign -action parameter values.
	 managefbvol command - Added tier and assignorder parameters. Added sfdisable, tierassign, and tierunassign -action parameter values.
	• chextpool command - Added tierunassign parameter.
	• showextpool command - Added tier parameter.

Version 7, Release 0, Modification 5

Function	Description
New commands	• mkipsec command - Creates an Internet Protocol Security (IPSec) connection on the DS8000 by importing a configuration file that contains a connection definition to the hardware management console (HMC).
	• chipsec command - Modifies an existing Internet Protocol Security (IPSec) connection.
	• rmipsec command - Deletes an Internet Protocol Security (IPSec) certificate from the hardware management console (HMC).
	• setipsec command - Allows you to manage Internet Protocol Security (IPSec) controls.
	• 1sipsec command - Displays a list of defined Internet Protocol Security (IPSec) connections.
	• mkipseccert command - Imports an Internet Protocol Security (IPSec) certificate to the DS8000.
	• 1sipseccert command - Displays a list of Internet Protocol Security (IPSec) certificates.
	• rmipseccert command - Deletes an Internet Protocol Security (IPSec) connection definition from the IPSec server.
Modified command	• offloadfile command - Added new parameters (-ipsec, -ipsecall, -ipsecraw, and -ipsecstat) to the syntax.
IPsec tasks	Added instructions for using DS CLI commands to perform IPSec tasks. For more information, see "Internet Protocol Security (IPSec) tasks" on page 511.

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

Version 7, Release 0

Function	Description
Modified commands	• mkuser command - Character limit enforced on username. Additional note added pertaining to enclosing password in quotes.
	• 1sflash command - Added new parameter (-dataset) to the syntax.
	• pausegmir command - Added new parameter (-withsecondary) to the syntax.
	• 1sgmir command - Added new Paused with Secondary Consistency and Paused because Resume Failed copy states.
	• showgmir command - Added new Paused with Secondary Consistency and Paused because Resume Failed copy states.
	• reverseflash command - Modified the default copy behavior when used with both the -tgtse and -fast parameters.
	• 1spprc command - Added new <i>Suspended</i> state reason codes to the report definitions.
	• mkfbvol command - Added new parameter (-t10dif) to the syntax.

Function	Description
Informational updates	 Documentation of Easy Tier default settings in the DS Storage Manager is updated. The default setting for I/O Priority Manager is now Disabled. Removed the Java 1.4.2 packages from the installation CD. These packages were previously included, although never required. You may acquire these packages from older DS CLI installation CDs at: ftp://ftp.software.ibm.com/storage/ds8000/updates/ DS8K_Customer_Download_Files/CLI/ The DS CLI installation CD previously contained the following Java packages: \IMAGES\HMC\IBMJava2-JRE-1.4.2-0.0.i386.rpm \IMAGES\HMC\ibm-java2-jre-142.exe \IMAGES\HMC\jre14-20040626.tar.gz

Version 6, Release 3, Modification 1

Function	Description
Modified commands	• The 1spprc command - Added new <i>Suspended</i> state reason codes to the report definitions.
	• The lsgmir and showgmir commands - Added new Paused with Secondary Consistency and Paused because Resume Failed copy states.
	• The pausegmir command - Added new parameter (-withsecondary) to the syntax.
	• The mkfbvol command - Added new parameter (-t10dif) to the syntax.
Removed commands	• Removed the following deprecated commands and older reference information from this publication. This information is now available in the <i>Archived CLI Information</i> section of the IBM System Storage DS8000 Information Center.
	 DS6000[™] remote support and notification commands (setplex, showplex, setdialhome, setsmtp, setsnmp, setsim, setcontactinfo, showcontactinfo, and testcallhome)
	- DS6000 PE package commands (offloadss, mkpe, sendss, sendpe, lsss, and lspe)
	 DS6000 problem log commands (closeproblem, and lsproblem)
	– setoutput command
	- The Command equivalents topic.
	- The Output field descriptions topic.
	 The Configuring the DS8000 (using DS CLI) for use with the Tivoli Storage Productivity Center Replication Manager topic.
	 Instructions for installing the DS CLI on an OpenVMS system.

Function	Description
Modified commands	• The help command - Add new parameter (-match) to the syntax.
	 The offloadfile command - Added new parameters (-sdocert, -auditlog, and -hmc) to the syntax.
	• The 1sfbvol command - Added a new field (Datatype FB 512t) to the report definitions.
	• The showfbvol command - Added a new variable (512t) for -datatype to the syntax diagram, and added a new field (Datatype FB 512t) to the report definitions.
	• The showfbvol command - Added new fields to the report definitions.
	• The showckdvol command - Added new fields to the report definitions.
	• The mksestg command - Changed the -vircap parameter from required to optional.
	• The lshostconnect command - Add new parameter (-wwpn) to the syntax.
	• The mkhostconnect command - Added the -wwpn parameter as an alias to -wwname .
	• The setenv command - Added new parameters (-locale, -timeout, -fullid, and -maxNumReports) to the syntax.
	• The showenv command - Added new fields to the report definitions.
	• The showextpool command - Added new fields to the report definitions.
	• The lsextpool command - Added new fields to the report definitions.
New commands	• The lsaccess command displays the access settings of a hardware management console (HMC).
	• The chaccess command allows you to change one or more access settings of a hardware management console (HMC). Only users with admin authority can access this command.

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

Version 6, Release 1

Function	Description
	Updated the instructions for installing the DS CLI application using the graphical mode with changes to the way the QDSCLI library output is stored.
Error code descriptions	Added error code descriptions to the installation instructions.

Function	Description
Modified commands	• mkckdvol command - Added a new parameter (-perfgrp) to the syntax.
	• chckdvol command - Added a new parameter (-perfgrp) to the syntax.
	• lsckdvol command - Added a new parameter (-perfgrp) to the syntax. Added a new field (perfgrp) to the report definitions.
	• showckdvol command - Added a new field (perfgrp) to the report definitions.
	• mkfbvol command - Added a new parameter (-perfgrp) to the syntax.
	• chfbvol command - Added a new parameter (-perfgrp) to the syntax.
	• 1sfbvol command - Added a new parameter (-perfgrp) to the syntax. Added a new field (perfgrp) to the report definitions.
	• showfbvol command - Added a new field (perfgrp) to the report definitions.
	• 1skey command - Added a new Activation Key (DS8000 I/O Priority) to the table.
	• chrank command - Added a new parameter (-quiet) to the syntax.
	• 1srank command - Added two new variables (<i>unassignedreserved</i> and <i>depopulationerr</i>) to the syntax. Added new fields (Unassigned Reserved, and Depopulation Error) to the report definitions.
	• showrank command - Added new fields (Depopulation Error and Unassigned Reserved) to the report definitions.
	• 1sfbvol command - Added a new variable (<i>managed</i>) to the syntax. Added a new field (managed) to the report definitions.
	• showfbvol command - Added a new field (managed) to the report definitions.
	• lsckdvol command - Added a new variable (<i>managed</i>) to the syntax. Added a new field (managed) to the report definitions.
	• showckdvol command - Added a new field (managed) to the report definitions.
	• mkckdvol command - Added a new parameter (-resgrp) to the syntax.
	• chckdvol command - Added a new parameter (-resgrp) to the syntax.
	• lsckdvol command - Added a new parameter (-resgrp) to the syntax. Added a new field (resgrp) to the report definitions.
	• showckdvol command - Added a new field (resgrp) to the report definitions.
	• mkfbvol command - Added a new parameter (-resgrp) to the syntax.
	• chfbvol command - Added a new parameter (-resgrp) to the syntax.
	• 1sfbvo1 command - Added a new parameter (-resgrp) to the syntax. Added a new field (resgrp) to the report definitions.
	• showfbvol command - Added a new field (resgrp) to the report definitions.
	• mklcu command - Added a new parameter (-resgrp) to the syntax.

Function	Description
Modified commands	• chlcu command - Added a new parameter (-resgrp) to the syntax.
	• 1s1cu command - Added a new parameter (-resgrp) to the syntax. Added a new field (resgrp) to the report definitions.
	 showl cu command - Added a new field (resgrp) to the report definitions.
	• chlss command - Added a new parameter (-resgrp) to the syntax.
	• IsIss command - Added a new parameter (-resgrp) to the syntax. Added a new field (resgrp) to the report definitions.
	• showlss command - Added a new field (resgrp) to the report definitions.
	• mkuser command - Added a new parameter (-scope) to the syntax.
	• chuser command - Added a new parameter (-scope) to the syntax.
	• Isuser command - Added new parameters (-s , -1 , and -scope) to the syntax. Added a new field (Scope) to the report definitions.
	• showuser command - Added a new field (Scope) to the report definitions.
	• setauthpol command - Added a new variable (<i>rmallmap</i>) and new parameters (-extgroup, -extuser, -dsgroup, and -scope) to the syntax.
	• showauthpol command - Added a new parameter (-revmap) to the syntax.
	• 1spprcpath command - Added a new field (PPRC CG) to the report definitions.
	• chpass command - Added new parameters (-age , -length , -history , and -reset) to the syntax.
	 showpass command - Added new fields (Password Expiration, Failed Logins Allowed, Password Age, Minimum Length, and Password History) to the report definitions.
	• showauthpol command - Added new fields (expire , age , fail , length , and history) to the report definitions.
	• showuser command - Added a new field (DaysToExpire) to the report definitions.
	• Isuser command - Added a new parameter (User_Name) to the syntax.
	• chsi command - Added a new parameter (-iopmmode) to the syntax.
	• showsi command - Added a new field (IOPMmode) to the report definitions.
	• diagsi command - Added a new parameter (-quiet) and a new variable (<i>saveuilogs</i>) to the syntax.
New commands	• The lsperfgrp command allows you to view a list of performance groups and information for each performance group in the list.
	• The lsperfgrprpt command allows you to view a list of performance reports for the given set of performance groups, or all if none are specified.
	• The lsperfrescrpt command allows you to view a list of performance reports for a given resource or set of resources of a given type.
	• The who command allows you to view authentication information for users who are currently logged in.
	• The chresgrp command allows you to change a resource group object on a storage image.
	• The lsresgrp command allows you to view a list of resource group objects on the storage image.
	• The manageresgrp command allows you to manage the contents of a resource group object on a storage image.
	• The mkresgrp command allows you to create a resource group object on a storage image.
	• The rmresgrp command allows you to remove a resource group object on a storage image.
	• The showresgrp command allows you to view the detailed properties of a resource group.

Function	Description
Modified commands	• chlss command - Added a new parameter (-ss) to the syntax.
	• 1sckdvol command - Added new values for <i>configstate</i> to the report definitions.
	• showckdvol command - Added new values for <i>configstate</i> to the report definitions. Added new fields (migrating and migratingfrom) to the report definitions.
	• 1sfbvol command - Added new values for <i>configstate</i> to the report definitions.
	• showfbvol command - Added new values for <i>configstate</i> to the report definitions. Added new fields (migrating and migratingfrom) to the report definitions.
	• showextpool command - Added new fields (%migrating(in) and %migrating(out)) to the report definitions.
	• chextpool command - Added new parameters (-merge and -quiet) to the syntax.
	• showrank command - Added new fields (migrating(in) and migrating(out)) to the report definitions.
	• 1skey command - Added new activation key, IBM Easy Tier.
	• chsi command - Added new parameters (-etautomode and -etmonitor) to the syntax.
	• showsi command - Added new fields (ETAutoMode and ETMonitor) to the report definitions.
	• managereckey command - Added new values (enable and disable) for the -action parameter to the syntax.
	• 1skeygrp command - Added new values (disabled, enableauthpend, and disableauthpend) for the -reckeystate parameter to the syntax, and also to the report definitions.
	• showkeygrp command - Added new fields (disabled, enableauthpend, and disableauthpend) to the report definitions.
	• rmckdvol command - Added a new parameter (-force) to the syntax.
	• rmfbvol command - Added a new parameter (-force) to the syntax.
	• showgmir command - Added a new parameter (-session) to the syntax. Added new fields (unowned and recovering) to the report definitions.
	• 1skeygrp command - Added new field (datakeycreated) to the report definitions.
	• showkeygrp command - Added new field (datakeycreated) to the report definitions.
New commands	• The manageckdvol command allows you to initiate a change on count key data (CKD) volumes.
	• The managefbvol command allows you to initiate a change on fixed block (FB) volumes.
	• The offloadfile command allows you to offload the Easy Tier performance data.
	• The lsgmir command displays a list of Global Mirror for the storage image of the specified logical subsystem.
	• The managekeygrp command allows you to manage an encryption key group.

This table provides the current technical changes and enhancements for both the DS8000 and DS6000 series. It combines the information from the former DS8000 and DS6000 Command-Line Interface User's Guides.

Function	Description
Modified commands	Added support for IBM DS8000 Encryption Recovery Key.
	• Added information about changes to the Windows PATH environment variable. See "Upgrading the DS CLI on your system" on page 16.
	• showsu command - Added new values for <i>Config</i> to the report definitions.
	• lskeygrp command - Added the <i>unconfigured</i> state and the reckeystate parameter to the syntax diagram. Added new fields (<i>reckeycreated</i> , <i>reckeystate</i> , and <i>unconfigured</i>) to the report definitions.
	• showkeygrp command - Added new fields (<i>reckeycreated</i> , <i>reckeystate</i> , and <i>unconfigured</i>) to the report definitions.
	• setauthpol command - Added secadmin authority.
	• mkuser command - Added secadmin authority.
	• chuser command - Added secadmin authority.
	• rmuser command - Added secadmin authority.
	• 1skeygrp command - Added <i>label2</i> to the report definitions.
	• mkkeygrp command - Added <i>label2</i> to the syntax diagram.
	• showkeygrp command - Added <i>label2</i> to the report definitions.
New commands	• The mkreckey command allows you to create a new encryption recovery key.
	• The managereckey command allows you to manage an existing encryption recovery key.
	• The rmreckey command allows you to remove an encryption recovery key.

Version 4, Release 3

Function	Description
Modified commands	Added support for IBM DS8000 Thin Provisioning.
	• mkckdvol command - Added a new value (ese) to the syntax for the -sam parameter.
	• lsckdvol command - Added a new value (ese) to the syntax for the -sam parameter. Also added ese as a value for the -sam parameter in the report definitions.
	• showckdvol command - Added new fields (realextents and virtualextents) and a new value for the -sam parameter to the report definitions.
	• mkfbvol command - Added a new value (ese) to the syntax for the -sam parameter.
	• 1sfbvol command - Added a new value (ese) to the syntax for the -sam parameter. Also added ese as a value for the -sam parameter to the report definitions.
	• showfbvol command - Added new fields (realextents and virtualextents) and a new value for the -sam parameter to the report definitions.
	 mksestg command - Added new parameters (-wait and extentpool_ID) to the syntax.
	 chsestg command - Added new parameters (-vircap capacity, -repcap capacity, -reppercent percentage, -quiet, and -wait) to the syntax.
	 showsestg command - Added new fields (reqvircap(GiB/Mod1) and reqrepcap(GiB/Mod1)) to the report definitions.
	• lskeygrp command - Added a new field (reckeystate) to the report definitions.
	• showkeygrp command - Added a new field (reckeystate) to the report definitions.
	 setauthpol command - Added a new option (secadmin) to the -action setmap, -action addmap, and -action rmmap parameters.
	• mkuser command - Added a new option (secadmin) to the -group parameter.
	• chuser command - Added a new option (secadmin) to the -group parameter.
	• The glossary has been removed from this publication. You can view the glossary in the DS8000 Information Center at the following Web site: IBM DS8000 series online product documentation (www.ibm.com/support/knowledgecenter/ST8NCA/product_welcome/ds8000_kcwelcome.html)
	• lsflash , lsremoteflash , and lspprc commands - Replaced the "allowTgtSE" field with "isTgtSE" in the report definitions.
	• setioport command - Removed the scsi-fcp/ficon value as a topology option.
	• showioport command - Added a new field (physloc) to the report definitions.
New commands	• The echo command allows you to specify whether or not the dscli will echo each specified command, or to display a user-specified string.
	• The mkreckey command allows you to create a new encryption recovery key.
	• The managereckey command allows you to manage an existing encryption recovery key.
	• The rmreckey command allows you to deconfigure an encryption recovery key.

Function	Description
Updated information	• Modified the way that the dscli sets a default devid. If the devid is not set using the profile file or using the setenv command, and if the hardware management console detects only one storage image, then the devid environment variable will be set to its storage image id.
	• The DS command-line interface installer, previously InstallShield Multi-Platform (ISMP), has been replaced by the InstallAnywhere installer.
	• As a convenience, along with the new merged DS CLI User's Guide, copies of both the DS6000 and DS8000 Message Reference documents are now included as part of the DS CLI installation.
	• Added new metrics that provide information about fibre-channel errors.
	• Added Windows Vista, Windows Server 2008, and AIX [®] 6.1 to the list of supported operating systems.
	• This publication now abbreviates Binary Gigabytes as the industry standard <i>GiB</i> (1 GiB=2^30B) instead of <i>GB</i> . All command input parameters and output displays remain unchanged except for some modifications in the Space Efficient Storage commands.
Modified commands	• mkflash command - Added a new parameter (-pmir) to the syntax.
	• resyncflash command - Added a new parameter (-pmir) to the syntax.
	• reverseflashcommand - Added a new parameter (-pmir) to the syntax.
	• 1sflash command - Added a new parameter (-pmir) to the syntax.
	• rmflash command - Added a new parameter (-cprm) to the syntax.
	• 1sremoteflash command - Added new information to the report (Timeout and Pmir), and replaced the <i>SS_ID</i> variable with <i>Source_LSS_SSID</i> .
	• mkuser command - Added a new parameter (-pol) to the syntax, which allows you to specify the name of a basic authentication policy.
	• chuser command - Added a new parameter (-pol) to the syntax, which allows you to change the name of a basic authentication policy.
	• rmuser command - Added a new parameter (-pol) to the syntax, which allows you to remove a user from a specific basic authentication policy.
	• Isuser command - Added a new parameter (-pol) to the syntax, which allows you to list all of the users in a specific basic authentication policy.
	• showuser command - Added a new parameter (-pol) to the syntax, which allows you to view detailed account information for all of the users in a specific basic authentication policy.
	• chpass command - Added a new parameter (-pol) to the syntax, which allows you to change the password for users in a specific basic authentication policy.
	• showpass command - Added a new parameter (-pol) to the syntax, which allows you to view the properties of passwords for users in a specific basic authentication policy.
	 Isarray, showarray, Isarraysite, showarraysite, mkrank, Isrank, showrank, mkextpool, Isextpool, showextpool, showsp, showsi, and Isddm commands - Added support for disk encryption.
	• dscli command - Added IPV6 support to the hardware management console connection.
	• mksestg and chsestg commands - Added capacity type <i>mod1</i> .
	• showarraysite , lsarraysite , lsddm , showarray , and lsarray commands - Added support for high capacity SATA disk drives.

Function	Description
Modified commands	• failbackpprc command - Added a flag <i>-tgtse</i> to specify that the PPRC secondary volume is a space efficient volume.
	• showfbvol and showckdvol commands - Added <i>zHPFRead</i> and <i>zHPFWrite</i> fields to the result of the commands when the -metrics option is specified.
	• chsestg and mksestg commands - Changed the parameter -captype default value to <i>mod1</i> when the storage is CKD storage.
	• Issestg and showsestg commands - Changed the capacity unit to <i>mod1</i> in the result for CKD storage.
	• mkckdvol and chckdvol commands - Added the parameter -captype.
	• showckdvol command - Added cap (Mod1) field to display the capacity in <i>mod1</i> units.
	• Added <i>echo</i> and <i>echoprefix</i> variables to the dscli profile.
New commands	• The mkauthpol command allows you to create an empty authentication policy.
	• The lsauthpol command displays a list of all the authentication policies on the storage image.
	• The rmauthpol command allows you to remove an authentication policy.
	• The chauthpol command changes the general attributes of an authentication policy, such as the policy name and the activation state.
	• The cpauthpol command copies an existing authentication policy to a new policy.
	• The showauthpol command displays detailed properties of a specified authentication policy.
	• The testauthpol command allows you to test a specified authentication policy.
	• The setauthpol command modifies policy attributes that apply to a specific type of authentication policy, changing the contents of the policy.
	• The whoami command displays authentication information for the current user.
	• The mkkeymgr command creates an entry for the key server on the storage complex.
	• The chkeymgr command updates the attributes of the key server entry on the storage complex.
	• The rmkeymgr command removes a key server entry on the storage complex.
	• The lskeymgr command displays a list of the key server entries that are on the storage complex.
	• The mkkeygrp command creates an entry for the key server encryption key group on the storage image.
	• The rmkeygrp command removes an entry for the key server encryption key group on a specified storage image.
	• The lskeygrp command displays a list of the key server encryption key group entries on the specified storage image.
	• The showkeygrp command displays detailed information for a specified key server encryption key group entry on the storage image.
	• The setenv command allows you to set the environment variables.
	• The showenv command displays the environment variables.

Function	Description
Modified commands	• lshostconnect command- Changed the display of the atchtopo and speed attributes in the lshostconnect command report.
	 showhostconnect command- Changed the display of the atchtopo and speed attributes in the showhostconnect command report.
	• showsestg command- The percentages for the repcapalloc and vircapalloc values are now displayed as whole numbers.
	• Issestg command- The percentages for the repcapalloc and vircapalloc values are now displayed as whole numbers.
	• mksestg command- Added a note to the -vircap and -reppercent parameters.
	 mkckdvol command- Added a new parameter (-datatype) to the syntax and new information to the -3380 and -cap parameters.
	 chckdvol command- Added a new parameter (-datatype) to the syntax and new information to the -cap parameter.
	• lsckdvol command- Added 3390–A sub parameter to the -datatype parameter in the syntax and added 3390–A to the report for deviceMTM. (For DS8000 models only)
	"RAID 6 overview" was added to Chapter 1.
	• mkarray command- Added a new parameter (-raidtype 6) to the syntax. (For DS8000 models only)
	• Isarray command- Added a new parameter (-raidtype 6) to the syntax and added RAID type 6 to the report. (For DS8000 models only)
	• showarray command- Added RAID type 6 to the report. (For DS8000 models only)
	• Isrank command- Added a new parameter (-raidtype 6) to the syntax and added RAID type 6 to the report. (For DS8000 models only)
	• showrank command- Added RAID type 6 to the report. (For DS8000 models only)
	• Ispprcpath command- Added "System Reserved Path" to the Failed Reason report and added the state "Degraded" to the report.

Chapter 1. Introduction to the DS8000 series

The IBM DS8000 series is a high-performance, high-availability, and high-capacity series of disk storage that support continuous operations. The DS8000 series is built on IBM POWER microprocessors in dual shared processor complexes.

Note:

The DS8000 version of the DS CLI can connect to the DS6000 series, and all appropriate commands work correctly. Commands that are specific to the DS6000 series are documented only in the Archived CLI Information section of the IBM DS8000 Information Center and in versions of the DS CLI User's Guide earlier than version GC27-4212-00.

Overview of the DS8000 series

The DS8000 series offers various choices of base and expansion models so that you can configure storage units that meet your performance and configuration needs:

DS8100

The DS8100 features a dual two-way processor complex and support for one expansion model.

DS8300

The DS8300 features a dual four-way processor complex and support for up to four expansion models.

DS8700

The DS8700 provides the option of a dual two-way processor complex or dual-four way processor complex. A dual four-way processor complex provides support for up to four expansion models. (Dual LPAR support is not available for the DS8700.)

DS8800

The DS8800 provides the option of a dual two-way processor complex or dual-four way processor complex. A dual four-way processor complex provides support for up to three expansion models.

DS8870

The DS8870 is equipped with IBM POWER7+[™] based controllers. It integrates high-performance flash enclosures and flash cards to provide a new and higher level of performance for the DS8870. The flash enclosures and flash cards are supported in the Enterprise Class, Business Class, and All Flash configurations. The DS8870 All-Flash configuration provides twice the I/O bays and up to twice the host adapters as the standard DS8870 single frame configuration. DS8870 continues to be available in a standard configuration with disk drives and flash drives in the Enterprise Class and Business Class configurations.

All DS8000 series models consist of a storage unit and one or two management consoles, two being the recommended configuration.

The DS8000 series offers a range of features including automated storage tier optimization, point-in-time copy functions with IBM FlashCopy[®], Remote Mirror and Copy functions with Metro Mirror, Global Copy, Global Mirror, Metro/Global Mirror, Multi-Target Metro Mirror, IBM z/OS[®] Global Mirror, and z/OS Metro/Global Mirror.

In addition to the DS Command-Line Interface (CLI), the following management capabilities help you manage your DS8000 functions:

• IBM DS8000 Storage Management GUI. The new DS8000 Storage Management GUI enables easier, more effective management of logical and system configurations for the DS8000.

- IBM DS Open application programming interface (API). The DS Open API can be used to automate configuration management through customer-written applications. The DS Open API presents another option for managing storage systems by complementing the use of the DS Storage Manager and the DS command-line interface.
- IBM Tivoli Storage Productivity Center. The IBM Tivoli Storage Productivity Center complements the DS Storage Manager by providing advanced capabilities that can help you centralize the management of your storage environment. With the Tivoli Storage Productivity Center, it is possible to manage and fully configure multiple DS8000 storage systems from a single point of control.
- IBM Tivoli Storage Productivity for Replication Manager. The Tivoli Storage Productivity for Replication Manager facilitates the use and management of Copy Services functions such as the remote mirror and copy functions (Metro Mirror and Global Mirror) and the point-in-time function (FlashCopy).

Note: For DS8000, you can have a maximum of 256 clients connected to a single storage unit server at the same time. Clients include all DS Storage Managers, DS command-line interfaces, DS open application interfaces, and IBM Tivoli Storage Productivity Center for Replication sessions. However, you must not simultaneously start more than 100 client sessions including DS CLI sessions. Starting more than 100 sessions simultaneously can result in connection problems.

To learn additional information about the DS8000 series:

- See the IBM DS8000 Introduction and Planning Guide for your model.
- View IBM Knowledge Center which is an information database that provides you with the opportunity to quickly familiarize yourself with the major aspects of the DS8000 series and to easily recognize the topics for which you might require more information. Because the information is all in one place rather than across multiple publications, you can access the information that you need more efficiently and effectively.

To view DS CLI commands reference and usage information in Knowledge Center, use the search or filtering functions, or find it in the navigation by clicking **System Storage > Disk systems> Enterprise Storage Servers> DS8000** and see Reference in the navigation. Go to the IBM Knowledge Center website to learn more.

• View the e-Learning modules that are available from IBM Knowledge Center website. The e-Learning modules provide animated presentations that describe the installation, configuration, management, and servicing of the DS8000 series.

Chapter 2. 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, the tasks that are involved in upgrading your system (particularly if your network includes IBM TotalStorage Enterprise Storage Servers machine type 2105), and the operational limitations that are associated with the DS CLI.

While the DS CLI installation CD comes with the release bundle documentation for the DS8000, the installation CD ISO image files are also available online. Each entry in the following table will take you to a listing of release bundles and their corresponding DS CLI version numbers. At the bottom of each listing is a link to the FTP site where you can download the installation CD ISO image files. You can also go directly to the FTP site at DS CLI Downloads (ftp://ftp.software.ibm.com/storage/ds8000/updates/DS8K_Customer_Download_Files/CLI/).

DS8000 Series	DS8000 Bundle Listing
DS8100/DS8300	DS8100/DS8300 Bundle (www.ibm.com/support/ docview.wss?uid=ssg1S4000641)
DS8700	DS8700 Bundle (www.ibm.com/support/ docview.wss?uid=ssg1S4000853)
DS8800	DS8800 Bundle (www.ibm.com/support/ docview.wss?uid=ssg1S4000983)
DS8870	http://www-01.ibm.com/support/ docview.wss?uid=ssg1S4001056

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 5.1, 5.2, 5.3, 6.1, 7.1
- HP-UX 11.0, 11iv1, 11iv2, 11iv3
- HP Tru64 UNIX version 5.1, 5.1A
- Linux, Red Hat Advanced Server [AS] 3.0, Enterprise Server [ES] 3.0, Red Hat Enterprise Linux [RHEL] 4, 5, 6 and 7
- Linux, SUSE 8, 9, SUSE Linux Enterprise Server (SLES) 8, 9, 10, 11
- VMware ESX v3.0.1 Console
- Novell NetWare 6.5
- IBM i 5.4, 6.1, and 7.1
- OpenVMS 7.3-1 (or newer, Alpha processor only)
- Oracle Solaris 7, 8, 9
- Microsoft Windows Server 2000, 2003, 2008, 2012, Windows Datacenter, Windows XP, Windows Vista, and Windows 7, 8

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.

DS CLI preinstallation information

The IBM 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.

Specific preinstallation concerns

Consider the following specific concerns as you prepare to install the DS CLI.

DS6000 preinstallation specifics

After installation and before you can use the DS CLI commands on a DS6000 machine type, be aware of the following requirements:

- Your management console must be equipped with the DS Storage Manager graphical user interface (GUI).
- The GUI was installed as a full management console installation.
- Your storage system must be configured. You must use the DS Storage Manager for this initial configuration. The configuration process includes the following tasks:
 - Selecting your storage complex
 - Assigning your storage unit to the storage complex
 - Designating network information for the storage system

DS8000 preinstallation specifics

There are no preinstallation concerns for the DS8000.

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.

Supported host systems	Installation file location
IBM AIX (5.1, 5.2, 5.3, 6.1, 7.1)	IMAGES\HMC\Disk1\InstData\AIX\NoVM\ dsclisetup.bin
Hewlett-Packard-UX (11.0, 11iv1, 11iv2, 11iv3.)	IMAGES\HMC\Disk1\InstData\HP-UX\NoVM\ dsclisetup.bin
Linux (Red Hat 3.0 Advanced Server [AS] and Enterprise Server [ES]), RHEL 4, RHEL 5 Note: See the additional instructions following this table.	IMAGES\HMC\Disk1\InstData\Linux\NoVM\ dsclisetup.bin
Linux, SUSE 8, 9, SUSE Linux Enterprise Server (SLES) 8, 9, 10, 11 Note: See the additional instructions following this table.	IMAGES\HMC\Disk1\InstData\Linux\NoVM\ dsclisetup.bin
Oracle Solaris (7, 8, 9)	IMAGES\HMC\Disk1\InstData\Solaris\NoVM\ dsclisetup.bin

Supported host systems	Installation file location
HP Tru64 (5.1, 5.1A)	IMAGES\HMC\Disk1\InstData\HP-UX\NoVM\ dsclisetup.bin
Novell NetWare 6.5	IMAGES\HMC\Disk1\InstData\Windows\NoVM\ dsclisetup.exe
IBM i 5.4, 6.1, and 7.1	IMAGES\HMC\Disk1\InstData\Windows\NoVM\ dsclisetup.exe
OpenVMS 7.3-1 (or newer, Alpha processor only)	See the Archived CLI Information section of the IBM DS8000 series online product documentation (www.ibm.com/support/knowledgecenter/ST8NCA/ product_welcome/ds8000_kcwelcome.html) for information on how to install, use, and remove the DS command-line interface in an OpenVMS environment.
VMware ESX v3.0.1 Console	IMAGES\HMC\Disk1\InstData\Linux\NoVM\ dsclisetup.bin
Microsoft Windows Server 2000, 2003, 2008, 2012, Windows Datacenter, Windows Windows Vista, and Windows 7, 8	IMAGES\HMC\Disk1\InstData\Windows\NoVM\ dsclisetup.exe

- v Java Platform (Java 8 or later) must be installed on your system. Before Release 7.3, the DS CLI only required Java 1.4.2+. However, starting with Release 7.3, the DS CLI requires Java 8 or later. 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).
- v For an AIX installation:
 - 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

- v For a Windows installation:
 - The User Access Control (UAC) settings for Windows Vista, Windows 7 and later, or Windows Server 2008 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.
- v For an IBM i model installation:

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:

- 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

HP-UX

/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

HP Tru64 UNIX

/opt/ibm/dscli

IBM i /ibm/dscli

VMware

/opt/ibm/dscli

Novell NetWare

SYS:\IBM\dscli

• Regardless of the operating system and DS series that you use, activate your license activation codes (part of the DS Storage Manager postinstallation instructions). Then, you can use the CLI commands that are associated with Copy Services functions.

Mounting the DS CLI installation CD

Before you can initiate the DS CLI installation, most systems require that you mount the DS CLI installation CD. These instructions describe how to mount the installation CD on each of the supported operating systems. The Windows operating system does not require that you mount the CD. Because the OpenVMS system has additional installation requirements, the installation instructions for the OpenVMS system are explained in a separate topic.

Complete the following steps to mount the DS CLI installation CD in preparation for the DS CLI installation:

- 1. Log on to your host system as a root user or administrator.
- 2. Insert the DS CLI product CD into the CD drive. If a window opens for the CD drive, close the window.
- **3**. Mount the CD drive using the **mount** command according to your system. You can mount your CD drive using the following examples:

AIX Create a directory for the CD-ROM: mkdir /cdrom -p

Create a file system for the CD-ROM: crfs -v cdrfs -p ro -d $cd\theta$ -m /cdrom

where *cd0* represents the CD-ROM drive.

Mount the CD-ROM file system: mount /cdrom

HP-UX

Mount the CD-ROM file system by using the path name for your environment:

ioscan -funC disk | more
mount /dev/rdisk/disk? /<cdrom>

Note: The device name /dev/rdisk/disk? is the default format for the HP 'agile' mode as defined in HP-UX 11iv3. In older HP-UX versions, or HP-UX 11iv3 in the legacy mode, the device name format is /dev/dsk/c?t?d?.

Linux Enter the following command on Red Hat systems:

mount /dev/cdrom

Oracle Solaris

Enter the following command:

mkdir /mnt
mount -F hsfs -r /dev/dsk/c0t6d0s2 /mnt

Note: The device name /dev/dsk/c0t6d0s2 is the default name for Oracle Solaris. The device name might be different on your system depending on your hardware configuration.

Windows

You are not required to mount the CD if you are using this operating system.

HP Tru64 UNIX

Issue the following command: mount -t cdfs -o noversion /dev/rznn /mnt

where *nn* represents the number of CD-ROM drives.

Novell NetWare

You are not required to mount the CD if you are using this operating system.

4. Navigate to your CD drive and proceed with either the unattended (silent), console, or graphic installation.

Installing the DS CLI using the graphical mode

Complete this task to install the DS CLI on your system using the graphical installation mode. Users of Windows, Novell NetWare, UNIX, Linux, and System i[®] systems can install the DS CLI using the graphical mode.

Consider the following requirements before you install the DS CLI:

v You must have a version of Java 8 or later that is installed on your system in a standard directory. The DS CLI installer checks the standard directories to determine if a version of Java 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 8 or later.

- v If the DS CLI has been previously 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.
- v If you are installing onto a Novell NetWare system, you are directed to provide the following information:
 - The directory where your Windows drive is mapped
 - The location of the JAVA_HOME environment variable. The JAVA_HOME environment variable is where the Java runtime environment file is installed. During the DS CLI installation process, you must specify the location of the Java Runtime Environment (JRE) file.
- v System i 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.

Notes:

- 1. From the command line, you can use the *-i* parameter to specify any user interface mode when installing the DS CLI: -i [swing | console | silent]. The default mode for installing Windows is *swing*. The default for UNIX and Linux is *console* mode. You do not have to specify the mode in the command unless you want to use something other than the default mode.
- 2. While in console mode, you can type back to return to the previous screen, or quit to exit the installation.

You can install the DS CLI using the graphical mode with the help of an installation wizard. Before you can use the DS CLI, some operating systems require that you restart your system after the installation is complete. You might also be required to open a new command prompt window to start a DS CLI session.

Note: After you install the new version of the DS CLI, your old DS CLI sessions might be unusable.

Complete the following steps to install the DS CLI using the graphical mode:

1. Start the setup file that is appropriate for your operating system. 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\WoVM\dsclisetup.exe*.

The Introduction window is displayed.

Initially (for all types of installation), 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 one of the following messages:

v 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.

v 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. The manner in which you respond to this message depends on your operating system and your installation environment settings. If the installation fails because the correct version of Java is not found, see "Correcting the Java Virtual Machine Not Found Error" on page 15.

- 2. Click **Next** on the Introduction window to continue or **Cancel** to exit the installation. When you click **Next** the License Agreement window is displayed.
- Select I accept the terms of the License Agreement and click Next to continue. Select I do not accept the terms of the License Agreement and Cancel to exit the installation.
 In Windows when were seent the concernent on d slick Next the ColectTranst Content windows in the concernent of the concernent of the ColectTranst Content of the concernent of the concernent

In Windows, when you accept the agreement and click **Next**, the SelectTarget System window is displayed (See Step 4. Otherwise proceed to step 5.

- 4. (Windows, Novell NetWare, or OS/400[®] installation) Select the target system (Windows, Novell NetWare, or OS/400) where you want the DS CLI installed, and then click **Next** to continue or **Cancel** to exit the installation.
 - a. Select Windows as your target system for all systems except Novell NetWare and OS/400. When you select Windows and click **Next**, the Choose Install Folder window is displayed.
 - b. When you select Novell NetWare and click **Next**, the Novell Location window is displayed. Go to Step 6 to continue the installation.
 - **c.** When you select OS/400 and click **Next**, the OS/400 System Information window is displayed. Go to Step 7 to continue the installation process.
- **5.** Verify that the directory name that is shown in the Choose Install Folder window is the directory where you want to install the DS CLI. If it is not the correct directory, enter the directory path in the input field. Click **Next** to continue the installation. Click **Cancel** to exit the installation.

When you click **Next** to continue the installation, the Pre-Installation Summary window is displayed. Go to Step 8 to continue the installation process.

Note: If you are installing onto a System i system, a window that asks for the directory where Java is installed on the i5/OS is displayed when you click **Next**. Go to Step 7 to continue the installation process.

6. (Novell NetWare installation) When you select Novell NetWare and click **Next**, the Choose Install Folder window is displayed. Click **Next** to continue. Complete the information on the Set Novell NetWare Configuration window. You are asked to supply the location where the Windows drive is mapped and where the Java home directory that contains the version of Java you want to use is located. Click **Next** to continue the installation. Click **Cancel** if you want to exit the installation.

When you click **Next** to continue the installation, the Pre-Installation Summary window is displayed. Go to Step 8 to continue the installation process.

- 7. (OS400 installation) On the OS/400 System Information window, confirm that any previous versions of the CLI have been uninstalled and then click **Next**. On the Enter Sign On Credentials window, enter the iSeries system name, user name, and password. Click **Next** to continue. The Pre-Installation Summary window is displayed. Go to Step 8 to continue the installation process.
- 8. (Pre-Installation Summary window) Verify that the displayed information is accurate. This window provides the location where the command-line interface will be installed and specifies how much space it will occupy on your drive. Click **Install** to continue or **Cancel** to exit the installation. You can change the installation directory by clicking the **Previous** button.

When you click **Install** to continue the installation process, the Installation progress window is displayed.

9. (Installation progress window) This window provides the progress bar that reflects the progress of the installation as the files are installed. After the installation is complete, click **Next** to continue or **Cancel** to exit the installation.

When you click **Next** to continue the installation process, the Important Information window is displayed.

10. (Important Information window) This window provides you with the opportunity to read the README file. Click **Done** to complete the installation.

Notes:

- 1. The DS CLI is installed in the following two places in i5/OS:
 - v IFS directory *IBM/DS_CLI*. This directory contains the profiles, .EXE files, Java .JAR files, readme files, and so on.
 - v 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 have programmed anything using the output, you should 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 /opt/ibm/dscli

HP-UX

/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

HP Tru64 UNIX

/opt/ibm/dscli

IBM i /ibm/dscli

Novell NetWare SYS:\IBM\dscli

Installing the DS CLI using the console mode

Complete this task to install the DS CLI on your system using the 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.

Consider the following before you complete the installation of the DS CLI:

- v You must have a version of Java 8 or later that is installed on your system in a standard directory. The DS CLI installer checks the standard directories to determine if a version of Java 8 or later exists on your system. If this version is not found in the standard directories, the installation fails. If the installation fails because the correct version of Java is not found, see "Correcting the Java Virtual Machine Not Found Error" on page 15.
- v If the DS CLI was installed on your client or host system in the past, you must end any active sessions of the DS CLI before you run the DS CLI installation CD.
- v If you are installing onto a Novell NetWare system, you are directed to provide the following information:
 - The directory where your Windows drive is mapped
 - The directory where the JAVA HOME is located

Before you can use the DS CLI, some operating systems require that you restart your system after the installation is complete. You might also be required to open a new command prompt window to start a DS CLI session.

Complete the following steps to install the DS CLI 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 command on the command line: dsclisetup<.exe | .bin> -i console. For example, for Windows, type: dsclisetup.exe -i console or, for Linux, type: dsclisetup.bin -i console. For an installation onto an OS/400 system from a Windows operating system, type: setupwin32console.exe -os400. The Introduction screen is displayed.

Notes:

- a. You can use the -*i* parameter to specify any user interface mode when installing the DS CLI: -i [swing | console | silent]. The default mode for installing Windows is swing. The default for UNIX and Linux is console mode. You do not have to specify the mode in the command unless you want to use something other than the default mode.
- b. While in console mode, you can type back to return to the previous screen, or quit to exit the installation.

```
Preparing CONSOLE Mode Installation...

IBM System Storage DS Command Line Interface(created with InstallAnywhere by Macrovision)

InstallAnywhere will guide you through the installation of IBM System Storage DS Command

Line Interface. It is strongly recommended that you quit all programs before continuing


with this installation. Respond to each prompt to proceed to the next step in the

installation. If you want to change something on a previous step, type 'back'.

You may cancel this installation at any time by typing 'quit'.

PRESS <ENTER> TO CONTINUE:
```

4. Press Enter to continue. The License agreement screen is displayed.



5. Type Y and press **Enter** to accept the terms of the license agreement. If you are running the installer on Windows, the Select Target System screen is displayed.

```
Select Target System
------
Please select the appropriate target system:
->1- Windows
2- Novell NetWare
ENTER THE NUMBER FOR YOUR CHOICE, OR PRESS ENTER TO ACCEPT THE DEFAULT:
```

6. If you are installing on a Windows operating system, type 1. If you are installing on a Novell NetWare operating system, type 2. Press **Enter** if you want to select the default operating system, or press **Enter** after you have typed in your selection. The Choose Install Folder screen is displayed.

7. Enter the path where you would like to install the DS CLI, or press **Enter** to accept the default location. The confirmation message below is displayed.

```
INSTALL FOLDER IS: C:\Program Files\IBM\dscli
IS THIS CORRECT? (Y/N):
```

- **8**. Type Y and press **Enter** to continue installing the DS CLI in the specified location. Type N and press **Enter** to change the location. Type Y to confirm that the location is correct, and press **Enter**.
 - a. If you are installing on Windows, the Pre-Installation Summary window is displayed.

```
Pre-Installation Summary
------
Please Review the Following Before Continuing:
Product Name:
    IBM System Storage DS Command Line Interface
Install Folder:
    C:\Program Files\IBM\dscli
Disk Space Information (for Installation Target):
    Required: 28,988,352 bytes
    Available: 80,796,971,008 bytes
PRESS ENTER TO CONTINUE:
```

b. If you are installing the CLI on a Novell NetWare system, the Set Novell NetWare Configuration window is displayed. Enter the Novell location where the Windows drive is installed and press Enter. The Set Novell NetWare Configuration screen is displayed. Enter the location of the JAVA directory that you would like to use and press Enter. The Pre-Installation Summary screen is displayed.

```
Set Novell NetWare Configuration

Please indicate the Novell location:

Novell location(volume:directory) (DEFAULT: SYS:):
```

9. Press Enter to begin the installation process.

10. Press Enter to continue. Important information about the README file is displayed.

```
Important Information
Please read the information below.
IBM(R) System Storage(R) DS Command Line Interface
for Microsoft(R) Windows 2000(R), Windows 2003(R)
Host Systems
README
 Contents
1.0 About this README file
1.1 Who should read this README file
1.2 Help contacts
2.0 Where to find more information
3.0 Contents of Windows CLI package
4.0 Notices
5.0 Trademarks and service marks
                         1.0 About this README file
This README file tells you where to find user
information about the IBM System Storage DS
Command Line Interface (CLI) User's Guide and lists
PRESS <ENTER> TO CONTINUE:
```

11. Press **Enter** to continue reading until you reach the end of the important information. When you finish reading the important information, the CLI installation is complete.

Installing the DS CLI using the unattended (silent) mode

Complete this task to install the DS CLI using the unattended (silent) mode.

Consider the following before you install the DS CLI:

- You must have installed a version of Java 8 or higher on your system in a standard directory. The DS CLI installer checks the standard directories to determine whether a version of Java 8 or higher exists on your system. If this version is not found in the standard directories, the installation fails. If the installation fails because the correct version of Java is not found, see "Correcting the Java Virtual Machine Not Found Error" on page 15.
- v If the DS CLI was installed on your host machine in the past, ensure that you end any active sessions of the DS CLI before you run the DS CLI installation CD.
- v Silent mode installation on OS/400 and Novell NetWare systems is not supported.

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.

Before you can use the DS CLI, some operating systems require that you restart your system after the installation is complete. Or, you might be required to open a new command prompt window to start a DS CLI session.

Note: After you install the new version of the DS CLI, your old DS CLI sessions might be unusable.

Complete the following steps to install the DS CLI using the unattended (silent) mode:

- 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\WoVM\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:

v 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.

v 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 15.

- 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"

Correcting the Java Virtual Machine Not Found Error

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

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:

 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:

v 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.

v 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 ensuring that Java 8 or later is installed, complete one of the following actions to correct the Java Virtual Machine Not Found error:

- v 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: Because there is a space in the Program Files directory name, you are required to add quotes around the file 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.

v (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.

v (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. Beginning with Release 4.2, 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 *or* one ESS CLI) are 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 may 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 execute the rest of your scripts 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 has completed. You can therefore 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 using the graphical (swing) mode, unattended (silent) mode, or console mode.

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

Uninstalling the DS CLI from your system using graphical mode

Complete this task to uninstall the DS CLI from your system when the DS CLI is installed on a Windows, Novell NetWare, or UNIX system.

Notes:

- 1. The following procedure applies to uninstalling the DS CLI. This procedure cannot be used to uninstall other versions of the CLI.
- 2. 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. Run the uninstaller to remove the DS CLI from your system.

Note: Refer to the User's Guide that is installed with your current DS CLI for removal instructions.

- a. 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.
- b. Open the _uninst directory (On Windows, C:\Program Files\IBM\dscli_uninst\uninstaller.exe. On UNIX or Linux systems, /opt/ibm/dscli/_uninst/uninstaller.
- c. You can use the -*i* parameter to specify any user interface mode when you uninstall the DS CLI:
 -i [swing | console | silent]. The default mode for uninstalling Windows is *swing*. The default for UNIX and Linux is *console* mode. You do not have to specify the mode in the command unless you want to use something other than the default mode.

Note: While in console mode, you can type back to return to the previous screen, or quit to exit the installation.

- d. From the command prompt, specify the -i swing parameter to open the uninstaller in graphical mode. The Uninstall IBM System Storage DS[®] Command Line Interface window is displayed. Click the **Uninstall** button to complete the uninstallation process, or **Cancel** to cancel the uninstallation.
- e. The Uninstall Complete window is displayed after the uninstallation process completes. Click **OK** to close the window.
- 2. Alternately, you can use the Add/Remove Programs facility of the Windows operating system to uninstall the DS CLI from your system. When you complete the uninstallation steps, restart your system to complete the uninstallation. Complete the following steps to uninstall the DS CLI using the graphical mode.
 - a. Navigate to the Windows Control Panel and open the Add/Remove program facility.
 - b. Scroll the list of currently installed programs and click the listing for IBM DS Command-Line Interface.
 - c. Click the **Change/Remove** button and the Welcome window for the Uninstaller is displayed.
 - d. Click **Next** to continue or click **Cancel** to exit the removal process. When you click **Next**, the Confirmation window is displayed that shows the directory from which the DS CLI program is uninstalled.
 - e. Click **Remove** to continue or **Cancel** to stop the removal and exit the uninstallation process. Click **Back** to return to the previous window. When you click **Remove**, the Uninstallation Progress window is displayed. When the uninstallation process is finished, the Finish window is displayed, which contains a statement about the success or failure of the uninstallation process. Click **Finish** to close.

If the uninstallation program does not remove some information from your system, the Restart window is displayed. You must restart so that previously locked files are released and automatically deleted.

- f. Close the Add/Remove Programs window.
- g. Restart your system, if required (now or later), to complete the removal process.

Uninstalling the DS CLI 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 OpenVMS, Novell NetWare, or System i i5/OS.

Complete the following steps to successfully uninstall the DS CLI.

Notes:

- 1. If you are using Windows or Novell, you use the Add/Remove Programs feature to uninstall the DS CLI.
- 2. This uninstall process works only with DS CLI. No other versions of CLI can be removed with this process.
- Locate the uninstaller file in the /_uninst folder. If you selected the default directory, you can find the _uninst folder using the /opt/ibm/dscli path. The uninstaller file name is uninstaller.exe for Windows, or uninstaller for other platforms.
- Type the following command at the command prompt: <install directory>/_uninst/uninstaller -i silent
- 3. Press the Enter key. All the associated CLI files are uninstalled.

Note: On Windows, you might need to restart your computer to complete the removal of files that were locked during the uninstallation process.

Uninstalling the DS CLI 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 use of an X display.

Complete the following steps to uninstall the DS CLI using the console mode:

- 1. Enter the following command at a command prompt: <install directory>/_uninst/uninstaller -i console
- 2. The Uninstall IBM System Storage DS Command Line Interface page is displayed.

```
Preparing CONSOLE Mode Installation...

IBM System Storage DS Command Line Interface (created with
InstallAnywhere by Macrovision)

Uninstall IBM System Storage DS Command Line Interface

About to uninstall...

IBM System Storage DS Command Line Interface

This will remove features installed by InstallAnywhere. It will not

remove files and folders created after the installation.

Hit "Enter" to uninstall..

PRESS Enter TO CONTINUE:
```

3. Press **Enter** to continue. The DS CLI is now uninstalled from your system. You might have to restart to complete the removal of files that were locked during the uninstallation process.

Uninstalling the DS CLI from a System i model

This section contains information to help you uninstall the DS CLI from a System i model.

Because the DS CLI is installed on a System 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 System 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 System 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 System i model by completing the following steps:
 - 1. Delete the library 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 System i model.

• Uninstall 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 System i model directly

Complete this task to uninstall the DS CLI through the direct use of your System i model.

You cannot use the conventional DS CLI removal methods that are used on other systems because the installation of the DS CLI on your System i model was done from a remote system. The remote installation does not allow the creation of an uninstaller that can be used directly by your System i model for the removal process. However, it is possible to use your System i model directly (bypassing the uninstaller) to remove the DS CLI.

You cannot use the uninstaller that was created for the DS CLI when you originally installed the DS CLI because it was created for the remote system that you used for the installation and not for the System i model.

You can complete this procedure at any time. However, it is common to complete this procedure when you want to uninstall the DS CLI from your system, but you do not intend to complete an associated upgrade of the DS CLI.

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

Complete the following steps to uninstall the DS CLI through the direct use of your System i model:

1. Issue the following command from your i5/OS application:

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

Substitute your uninstall directory for yourdir.

2. Wait until the uninstall process is complete before you continue with your next process.

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

Complete this task to uninstall the DS CLI from your System i model using the remote method.

You cannot use the conventional DS CLI removal methods that are used on other systems because the installation of the DS CLI on your System i model has been done from a remote system. The remote installation does not allow the creation of an uninstaller that can be used directly by your System i model for the removal process. However, it is possible to use the remote removal method on your System i model to remove the DS CLI.

Ensure that the remote system that you use to uninstall the DS CLI is network-attached to the System i model and is a supported platform for DS CLI.

You can use the following remote method to uninstall the DS CLI from a System i model. You can complete this procedure at any time. However, it is common to complete this procedure when you want to upgrade the DS CLI, because the remote system that you are using to remove the DS CLI is typically the same system that you use for the upgrade.

To remove the DS CLI from your System i model using the remote method, complete the following steps:

- 1. Navigate to the _uninst folder on the Windows desktop. The _uninst folder was created when the DS CLI installer was initially run on the remote Windows machine.
- 2. Run the uninstaller.exe file to uninstall the DS CLI on the remote System i system. Enter the system name, user name, and password to complete the uninstallation.

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 or DS6000.

You must install the DS CLI before you complete these post-installation tasks. See Chapter 3, "Running the DS CLI," on page 35 for information about running the DS CLI, obtaining help, and interpreting exit codes.

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 or DS6000 models. 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.

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

- 1. Set your DS CLI default configuration settings. If this is a new DS8000 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. Enable the remote support and call home functions for your DS6000 model. You must provide customer contact and network information to enable these functions.
- **6**. Use the DS CLI to enable SNMP traps for Copy Services events and storage complex events on your storage unit.
- 7. Register for the My Support service for your DS6000 model.
- 8. Optionally, configure encryption on your encryption-capable DS8000 storage unit.
- 9. 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_name> must also include the full path name. If the -cfg profile 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.

OpenVMS system

For an OpenVMS operating system, the property value defaults to the logical name **SYS\$LOGIN**. 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: <i>string</i>	Specifies your user name for entering DS CLI commands. This variable is equivalent to option -user .				
password: <i>string</i>	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: <i>passwordfile</i>	Specifies a password file containing your password as an alternative to the variable <i>password</i> . This variable is equivalent to option -pwfile .				
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:	Specifies the command prefix to print before a command is executed.				
<i>prefix</i> none	• 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. 				

Table 6.	Profile	variables	(continued)
----------	---------	-----------	-------------

Variable	Description
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: <i>string</i>	Specifies the secondary Storage Manager IP address. This variable is equivalent to the command option -hmc2 . The command option -hmc2 overrides this default value.
locale: <i>code</i>	Specifies the language for the output on the local computer.

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. 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 s the number of lines per page with the rows variable. This variable is equivalen command option -p . The command option -p overrides this default value.						

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: <i>number</i>	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.				

Table 6. Profile variables (continued)

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 format
#
   xml
   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 using the DS CLI

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

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 allowed by the assignment are as follows:

admin (Administrator)

All users that you assign to the storage administrator user role have access to all DS6000 storage management console-server service methods, and all DS8000 storage image resources except those that are reserved for security administrator users.

secadmin (Security Administrator)

All users that you assign to the security administrator user role may initiate recovery key operations, and add other users to this role. Users in this role may not be assigned to any other user role, and users in any other user role may not 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.

service (Service Operator)

The service operator user role includes monitor authority, plus access to all management console-server service methods and resources, such as performing code loads and retrieving problem logs.

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.

Capability	Admin- istrator	Security Admin- istrator	Physical Operator	Logical Operator	Copy Services Operator	Mon- itor	Service Operator (DS6000 only)	No Access
Security user account management		X						
Recovery key management		X						
User account management	x							
Access audit log	X							
Update storage complex	x		x					
Power on/off storage image	х		x					
Update storage unit	Х		Х					
Update storage image	х		x					
Warmstart storage image	X		X					
Manage arrays, ranks, extent pools	X		x					
I/O port configuration	X		x					
Configuration recovery services (unfence volumes, discard pinned tracks, repair ranks,)	х		x					
Host configuration	Х		Х	Х				
Logical subsystem configuration	х		x	X				

Table 7. User role capabilities

Table 7. User ro	le capabilities	(continued)
------------------	-----------------	-------------

Capability	Admin- istrator	Security Admin- istrator	Physical Operator	Logical Operator	Copy Services Operator	Mon- itor	Service Operator (DS6000 only)	No Access
Volume configuration	x		х	x				
Add or remove volume group	x		Х	x				
Assign or unassign volume group to host connection	x		x	X				
Add or remove volumes to volume group	x		x	X				
Manage Copy Services (FlashCopy, PPRC, Global Mirror)	X		x		x			
Set Copy Services timeout values	x		Х		x			
Update user account password	x		Х	x	x	x	X	
Query FRUs and enclosures	x		Х	x	x	x	X	
Query configuration	x		X	x	x	x	x	
Query Copy Services	x		Х	x	x	x	X	
FRU management	Х		X				Х	
Problem management	x		Х				X	
Validate communication paths	x		x				X	
Activate code load	Х		X				Х	
Create a new PE package	x		х				x	
Manage storage unit IP addresses	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 log on. Users must change the password at the time of their initial log on. 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.

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.

- 1. Log in to the DS CLI in interactive command mode.
- Issue 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 using the DS CLI

Use the steps described in this task to activate your license activation codes. These codes must be activated before any configuration can be applied to your DS8000 or DS6000 network.

Some license activation codes include:

- Operating environment license for each storage unit that you own. (This license must be activated.)
- Copy Services, which can consist of the following features:
 - FlashCopy
 - Remote mirror and copy
- Parallel access volumes (PAV)
- IBM HyperPAV (DS8000 only)
- IBM FlashCopy SE (DS8000 only)

There are multiple codes that are associated with these features. For a complete list of activation codes, please 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.

Download your 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 2244 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 2244 serial number (which you can find on the License Function Authorization document).
- 2. For DS6000, in most situations, the DSFA application can locate your order confirmation code (OCC) when you enter the DS6000 (1750) serial number and signature. However, if the OCC is not attached to the 1750 record, you must assign it to the 1750 record in the DSFA application. In this situation, you must have the OCC (which you can find on the License Function Authorization document).

The DS CLI **applykey** command activates the licenses for your storage unit. The DS CLI **1skey** command verifies which type of licensed features are activated for your storage unit.

Complete the following steps to activate your license activation codes:

1. Log in to the DS CLI in interactive command mode.

- 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 has completed, the following message is displayed: Licensed Machine Code key xxxx, key xxxx successfully applied.
- 4. Verify that the keys have been activated for your storage unit by issuing the DS CLI **1skey** command as follows: 1skey -dev IBM.2107-75FA120

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

5. Press Enter and the following type of report is displayed:

OpenVMS system integration

You can adjust your OpenVMS system to obtain greater benefits from the use of the DS CLI. The hints and tips that are provided in this section show how to obtain these benefits through the optimal integration of the DS CLI into your OpenVMS system.

The following list provides the areas that you might consider for optimizing the use of the DS CLI in your OpenVMS system:

- Command Console LUN (CCL)
- OpenVMS system messages
- Message help
- Java Run Time Environment (JRE)
- Quota recommendations

Enhancing the command console LUN for DS CLI use

The OpenVMS operating system considers a Fibre Channel device with LUN ID 0 as Command Console LUN (CCL). These devices do not normally display when you issue the DS CLI **lshostvol** command. However, with adjustments, these devices can be displayed when you issue the **lshostvol** command. The following description provides the information that you need to make this enhancement work on your OpenVMS system.

Fibre Channel CCL devices have the OpenVMS device type GG, which result in OpenVMS device names in the form \$1\$GGA*n*. In contrast, Fibre Channel disk devices have the OpenVMS device type DG, which result in device names in the form \$1\$DGA*n*. Therefore, LUN 0 devices on OpenVMS are a special device type, different from disk devices.

The DS CLI **lshostvol** command displays the mapping of host device names or volume names to machine type 2105, 2107/242x, and 1750 volume IDs. That implies that all host devices belonging to

2105/2107/242x/1750 volumes are displayed. Therefore, CCL devices \$1\$GGA*n* are included in the lshostvol output for multiplatform consistency and to match the output of other DS CLI commands.

However, the inclusion of CCL devices can be confusing for users who expect that the **lshostvol** command displays only the disk devices. You can use the OpenVMS logical name **IBMDSCLI\$SHOW_GG_DEVICES** to modify the DS CLI behavior: If this logical name translates to an expression which evaluates as True in OpenVMS conventions (1, Y, YES, T, or TRUE), then the \$1\$GGA*n* CCL devices are shown in the command output. Otherwise, the \$1\$GGA*n* CCL devices are not shown.

The startup procedure IBMDSCLI\$STARTUP.COM defines the logical name

IBMDSCLI\$SHOW_GG_DEVICES as Y. If you want to suppress \$1\$GGA*n* CCL devices in the lshostvol command output, you can redefine the logical name after the startup procedure has been processed.

Enhancing the OpenVMS system messages

When you use the DS CLI, the application provides messages regarding the application processes, status, and errors. You also receive the OpenVMS system messages but they are displayed in a different format. You can make this situation less confusing by making the following adjustments.

The DS CLI messages are presented in an operating-system independent format. In contrast, native OpenVMS programs provide messages using the system message facility as displayed in the following format: **%facility-level-identification, text**.

To ensure that the OpenVMS command **SET MESSAGE** and customer-written tools that scan for such messages work correctly, the DS CLI provides each message using OpenVMS system services in addition to the operating system independent output. After displaying the OpenVMS message, the normal DS CLI message is provided unchanged. This ensures that the DS CLI messages are identical across platforms and that you can work with the DS CLI documentation.

However, these redundant messages can be confusing for users who are not familiar with OpenVMS. You can use the OpenVMS logical name **IBMDSCLI\$OPENVMS_MESSAGES** to modify the DS CLI behavior: If this logical name translates to an expression which evaluates as True in OpenVMS conventions (1, Y, YES, T, or TRUE), then the additional OpenVMS-formatted messages are presented. Otherwise, only the operating system independent DS CLI messages are shown.

The startup procedure IBMDSCLI\$STARTUP.COM defines the logical name

IBMDSCLI\$OPENVMS_MESSAGES as Y. If you want to suppress the OpenVMS-formatted messages, you can redefine the logical name after the startup procedure has been processed.

Enabling OpenVMS to use the DS CLI help

The DS CLI installation process offers the option to add modules to the system help library. If you enable OpenVMS with this option, you can use the DS CLI help.

The DS CLI installation process offers the option to add modules to the system help library SYS\$COMMON: [SYSHLP]HELPLIB.HLB and the system messages database SYS\$COMMON: [SYSHLP]MSGHLP\$LIBRARY.MSGHLP\$DATA. If you choose this option, the module IBMDSCLI is added as the top-level key to the help library, and the DS CLI status messages can be accessed using the HELP/MESSAGE/FACILITY=IBMDSCLI command. Additionally, the login procedure IBMDSCLI\$MANAGER:IBMDSCLI\$LOGIN.COM activates the message section file IBMDSCLI\$SYSTEM:IBMDSCLI_Messages_Shr.exe for the current process.

In every case, the installation process provides the following files in the directory which is referred by the logical name IBMDSCLI\$HELP:

IBMDSCLI_Ovr.hlp

A help library containing one module with the top-level key IBMDSCLI. You can add this library to the search list for help libraries in your OpenVMS system by defining appropriate logical names **HLP\$LIBRARY**, **HLP\$LIBRARY_1**, **HLP\$LIBRARY_2**, and so on.

IBMDSCLI_Messages.msghlp\$data

A message help data file with messages for facility IBMDSCLI. You can add this data file to the searchlist for message help files in your OpenVMS system by defining the logical name MSGHLP\$LIBRARY accordingly.

If you do not want the installation process to modify the OpenVMS system libraries, you can use these OpenVMS default logical names to integrate the DS CLI help information manually.

Java Runtime Environment considerations for DS CLI

The DS CLI login procedure **IBMDSCLI\$MANAGER:IBMDSCLI\$LOGIN.COM** defines **JAVA\$CLASSPATH** in the OpenVMS process logical name table and it overrides any existing Java classpath definition. If you want to use other Java-based software in the same process, you must redefine **JAVA\$CLASSPATH** so that it provides the classpath as a JAVA command parameter.

The following information provides an overview of how the installation of the DS CLI affects the Java environment of your OpenVMS system.

Because the DS CLI relies on Java Run Time Environment (JRE) V1.4.2, mandatory JRE files are installed in the directory tree that is referenced by the logical name **IBMDSCLI\$JRE**. This setup is according to HP guidelines. The login procedure **IBMDSCLI\$MANAGER:IBMDSCLI\$LOGIN.COM** calls the JRE setup procedure which defines several logical names and DCL symbols for usage by the Fast Virtual Machine.

If your OpenVMS host system uses other software that requires JRE but cannot run with the same JRE version as the DS CLI, users of that software can switch between different Java versions. To use different JRE versions, you must run a command procedure to set up the Java environment definitions for the version that you want to use in the given process. See the OpenVMS Java documentation at:

http://h18012.www1.hp.com/java/documentation /1.4.2/ovms/docs/user_guide.html

The DS CLI application-specific Java classes are bundled in Java Archive (.JAR) files in the directory referenced by logical name **IBMDSCLI\$LIBRARY**. These files must be included in the Java classpath. On OpenVMS, two logical names define the classpath:

CLASSPATH

For UNIX-style names. You can use a string inside single quotation marks that consists of colon-separated path names.

JAVA\$CLASSPATH

For OpenVMS specification syntax. You can specify multiple paths with a comma-separated expression (not enclosed in single quotation marks) as OpenVMS logical name search list. **JAVA\$CLASSPATH** overrides **CLASSPATH**, if **JAVA\$CLASSPATH** is defined.

Because of this override process, you might have to redefine the **JAVA\$CLASSPATH** to provide the class path as a JAVA command parameter. However, this JAVA command parameter is only required if you want to use other Java-based software in the same process.

Quota considerations for DS CLI

The JRE was designed to perform optimally on UNIX systems, where each process is given large quotas by default. On OpenVMS, the default behavior gives each process lower quotas so that many processes can co-exist on a system.

To get the best Java performance on OpenVMS, HP recommends that you set process quotas to match a typical UNIX system. HP also recommends these as minimum quota settings (except where noted). See these recommendations at the SDK for OpenVMS (h18012.www1.hp.com/java/documentation/1.4.2/ ovms/docs/user_guide.html)

To check whether your current process quotas fulfill the recommendations, you can run the following process: IBMDSCLI\$JRE:[LIB]Java\$Check_Environment.com.

Chapter 3. Running the DS CLI

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

Note: For the DS6000 model, you must ensure that you have installed the DS Storage Manager using the full-management console installation and that you have configured your domain. Without this domain configuration (which is a one-time process), you cannot use the DS CLI.

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 XP 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 into the DS CLI

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

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

HP-UX

/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

HP Tru64 UNIX

/opt/ibm/dscli

IBM i /ibm/dscli

Novell NetWare

SYS:\IBM\dscli

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 have 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*.

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, either in the directory where you installed the DS CLI or using the full path name for the DS CLI, you can specify this information using either of the following three methods:

- You may log in to the DS CLI without specifying any of this information on the command line and the application will prompt you to enter the information interactively. For example: /opt/ibm/dscli/dscli.
- You may log in to the DS CLI by specifying this information on the command line. For example: dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t.
- You may 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 into 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 that are provided in this guide 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.

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:

 Use the following command format to start a DS CLI session (Windows operating system): dscli -hmc1 9.1.23.456 -user admin -passwd my_password lssi -s -fullid -hdr off Here is an example of this same command in IBM i without the report delimiters: DSCLI SCRIPT(*NONE) HMC1('9.1.23.456') 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 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* will be 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.

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.

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

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

dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd password -script $^{\prime\prime}/bin/mkpprcpairs$

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.

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

(mknnrc	-dev	IBM 2107-1303561	-remoteday	IBM 2107_7504401	-type mmir 1000-103F:2300-233F
						-type qcp 1100-113F:2340-237F
						51 51
						-type mmir 1800-187F:2800-287F
						-type gcp 1200-127F:2500-257F
						-type mmir 1040-1054:2700-2714
						-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.

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 have used. It also allows you to 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:

- 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. The **setenv** command allows you to set or display command output format options. 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.

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 have 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 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.

Use the **1ssi** 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. This can save you input time when you have to process many transactions that cannot be part of a script.

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

- 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:

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></string>	Displays a list of commands that contain the specified string. This parameter cannot be used with any other parameters.
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.

Table 8. Parameters for the help command

Notes:

- You cannot use the **-s** and **-1** 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.

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.

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.

Table 9. DS CLI exit codes

Table 9. DS CLI exit codes (continued)

Code	Category	Description
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.

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.

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).

Note: 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.
- Beginning with Version 6 Release 1, 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. This section also provides information about how the DS CLI manages OpenVMS messages.

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 4. CLI commands

Use the command-line interface (CLI) commands for your configuration and storage management tasks.

About CLI commands

This topic provides 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. This means that 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 separated by a white space and not a comma (unlike flag parameters that allow multiple values). Some commands, like **1suser**, do not require a command parameter, because a default value of *all* is always assumed. For some commands, like **1sarraysite**, 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 provided in the command parameter string.

In the following example, **1srank** is the command name. **-dev** and **-1** 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.

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.

▶∢

-


*4

Keyword

► dscli

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

▶ — 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, [$a \mid b$] indicates that you can choose a, b, or nothing. Similarly, { $a \mid b$ } 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 <i>-p on</i> or <i>-p off</i> to turn the paging on or off; but if you specify only <i>-p</i> , 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 <i>char</i>	Sets the output to delimited output and the delimiter to the single character <i>char</i> . You must enclose <i>char</i> in single or double quotation marks if the character is a shell metacharacter (such as * or \t). If <i>char</i> is not specified, the CLI program returns a syntax error. A blank space, even when it is enclosed within quotation marks, is not a vali character as a delimiter. Note: You can use this option only with list (for example, Isuser, Iskey, Isserver) commands		

Parameters	Description
-hdr on off	Turns the header on or off. The default is on. You can specify <i>-hdr on</i> or <i>-hdr off</i> to turn the header on or off; but if you specify only <i>-hdr</i> , you get an error.
-bnr on off	Turns the banner on or off. The default is on. You can specify <i>-bnr on</i> or <i>-bnr off</i> to turn the banner on or off; but if you specify only <i>-bnr</i> , you get an error.
-v on ∣ off	Turns verbose mode on or off. The default is off. You can specify <i>-v</i> on or <i>-v</i> off to turn verbose mode on or off; but if you specify only <i>-v</i> , 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 -1 parameters together.
-1	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 -1 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 -1 parameters together.

Framework commands

This section contains the user interface framework commands for the DS command-line interface.

The following framework commands are available:

- **dscli** Starts the DS CLI. Use this command to perform storage management tasks from the command line.
- **echo** 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.
- 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 Allows you to set 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 are 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 21 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. If the second attempt also fails, the DS CLI attempts to connect to port 1718 with the legacy certificate used by ESS 2105 systems. This default behavior means that the DS CLI connects to any ESS 2105 or DS8000 system. However, checking multiple ports can cause a connection delay when a DS CLI connects to a DS8000 system or ESS 2105 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.

- 1718 Connect by using port 1718 only (ESS 2105 with legacy certificate).
- 1750 Connect by using port 1750 only (DS8000 before Release 7.2 with legacy certificate).
- **1751** Connect by using port 1751 only (DS8000 Release 7.2 and later 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.

► — exit-

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.

-1

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

"doetry herp -5 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.

```
▶ helpmsg—message_ID-
```

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 IBM 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.

xm1

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:

уууу	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.

xm1

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="unit8">5</DISPLAY>
<VALUE TYPE="string">>5</VALUE><PROPERTY>
<PROPERTY NAME="extpool_id"><DISPLAY TYPE="string">-
</DISPLAY><VALUE TYPE=string"></PROPERTY>
<PROPERTY NAME="extpool_id"><DISPLAY TYPE="string">-
</DISPLAY><VALUE TYPE=string"></PROPERTY>
<PROPERTY NAME="stgtype"></PROPERTY>
<PROPERTY NAME="stgtype"></PROPERTY></PROPERTY>
<PROPERTY NAME="stgtype"></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERTY></PROPERT
```

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.

► showenv-

Example

Enter the **showenv** command.

dscli> showenv

The resulting output

off paging rows 24 default format delim header on verbose off banner on IBM.2107-1301441 devid remotedevid IBM.2107-1301441 off echo echoprefix zh locale 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).

►► ver

Parameters

-s

(Optional) Specifies that the system displays the version of the command-line interface program. You cannot use the **-s** and **-1** parameters together.

*4

-1

(Optional) Specifies that the system displays the following version information:

- The version of the DS command-line interface
- The version of the DS Storage Manager (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 **-1** 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).

-1mc

(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	rageManager 7.7.7.0.20130628.1	
HMC DSCLI	7.7.30.180	
	=====Version===	
Storage Image	LMC	Bundle Version
IBM.2107-75YZ88	1 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 **-1** parameter is specified.

Security commands

This section contains commands that are used to maintain security, including user accounts and security, Internet Protocol Security (IPSec), and private network security.

User account and security commands

This section contains commands that are used to maintain user accounts and security.

The DS8000 system comes with an internal authentication and authorization service, called the Basic authentication service. This service also provides user management and is identical to the service provided in previous DS8000 models. The DS8000 system uses an external authentication service, such as a LDAP server, but still uses the internal authorization service to grant access to resources as defined by the DS8000 user group roles. These roles include administrator, storage operator, volume operator, Copy Services operator, service operator, and monitor.

An authentication policy defines the authentication service to be used. It also defines any parameters that are required to connect to and use that authentication service, and any mappings from that service's user groups to the DS8000 user group roles. It is assumed that any external authentication service provides tools for user account management, therefore the DS8000 system manages user accounts that are defined in the Basic authentication service. Currently, the DS8000 system defines two policy types, Basic and SAS (Storage Authentication Service). The Basic policy type uses the DS8000 Basic authentication service, and the SAS policy type uses the authentication service that is provided by an IBM SSPC (IBM System Storage Productivity Center).

With the introduction of the encryption recovery key, a "dual control" security process is required to protect the authorized usage of the recovery key. This dual control process requires two separate user accounts to process most recovery commands. If these accounts are owned by two different people, then the recovery key cannot be used by any one person to gain access to encrypted data. The first user role is the same 'admin' user role that has existed in previous versions of the DS8000 system, but it is now referred to as a storage administrator. The second user role, 'secadmin', is referred to as a security administrator. In the dual control process, the security administrator initiates a recovery key operation, and the storage administrator authorizes the operation that completes, or activates, the operation. The storage administrator cannot initiate any recovery key operation and a security administrator cannot authorize any recovery key operation.

A user that has security administrator authority cannot have the authority of any other user role. A user with any other user role cannot have the security administrator authority at the same time. This isolation of the security administrator authority is extended into the creation of a SAS policy by restricting who can specify mapping of external users and groups to the security administrator role to only those users with security administrator authority. This last restriction can create problems if the active SAS policy does not have any valid users or groups that are mapped to the security administrator role because no user can initiate any of the recovery key operations. To avoid these problems, the following SAS creation procedure is recommended:

- 1. Either the storage administrator or security administrator creates a new, or copies an existing SAS policy.
- 2. The security administrator specifies the mappings to the security administrator role.
- 3. The storage administrator specifies all other attributes and mappings of the SAS policy.
- 4. Either the storage administrator or security administrator tests the new policy.
- 5. The storage administrator activates the new policy.

The following user account and security commands are available:

chauthpol

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.

- **chpass** Changes the password expiration time, number of login attempts, minimum password age, minimum password length, and password history for a storage complex.
- **chuser** Modifies and locks or unlocks a user account. Users who do not have administrator authority, use this command to change an expired password, and create a password that is not known to the administrator who created their account.

cpauthpol

Copies an existing authentication policy to a new policy. You only can copy a policy within the same storage complex. Copying between different storage complexes is not supported.

lsauthpol

Displays a list of all the authentication policies on the storage complex.

Isuser Generates a report that lists the user names and access authority levels of the storage image user account.

managepwfile

Creates a password file for an existing ESS or DS user account. This command processes the password requirements for 2105, 2107/242x, and 1750 systems.

mkauthpol

Creates an empty authentication policy. Use the **setauthpol** command to set specific policy attributes.

mkuser Creates a DS CLI or a DS Storage Manager Basic authentication policy user account.

rmauthpol

Removes an authentication policy.

rmuser Removes a storage image user account. CLI users with administrative authority use this command to delete a user account file.

setauthpol

Modifies policy attributes that apply to a specific type of authentication policy, changing the contents of the policy.

showauthpol

Displays detailed properties of a specified authentication policy.

showpass

Generates a report that lists the number of days until the password expires, number of failed login attempts, password age, minimum password length, and password history that are associated with a password.

showuser

Generates a report that displays the details for an individual storage image user account.

testauthpol

Tests a specified authentication policy.

who Displays authentication information for the users who are currently logged in.

whoami Displays authentication information for the current user.

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. This command is not supported on DS6000 models.

▶▶	- chauthpol	
▶	pw— password—	

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 test2ibm 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.

>> —	npass	
►	nistory— numberresetpolname	-▶◀

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 16 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 DS Storage Manager basic user account.

All users can use this command to change their own password. However, users with administrator authority can also use this command to update any user's account password to modify user group authority, or to lock or unlock a user account. 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 16 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), NetWare, or OpenVMS 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)
- 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) 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.

Example

Invoking the chuser command

dscli> chuser -pw xy0abcde testuser

The resulting output

User tester 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. This command is not supported on DS6000 models.

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.

Isauthpol

The **lsauthpol** command displays a list of all the authentication policies on the storage complex. This command is not supported on DS6000 models.



Parameters

- S

(Optional) Displays only the policy name. You cannot use the **-s** and the **-1** parameters together.

-1

(Optional) Displays the default output and location. You cannot use the -s and the -1 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

NameTypeStateLocationPol1SASinactivehttp://yoursite.comPol2SASactivehttp://yoursite.commy policy1Basicinactive9.12.133.155

Report field definitions

Name*

Indicates the name of the authentication policy.

Туре

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 **-1** parameter is specified.

Isuser

The lsuser command returns a list of basic user account names and access authority levels.

► lsuser _______ -s______ -pol___pol__name_______-scope____user_resource_scope_____User_Name [...]_____

Parameters

- S

(Optional) Displays only the user names. You cannot use the -1 and the -s parameters together.

-1

(Optional) Displays the default output. You cannot use the **-1** 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	Scope
Testuser	service,op_copy_services	active	Product_A
Biguser	admin	active	*
Smalluser	op_storage	locked	Product_A

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)
- op_storage (Physical Operator)
- op_volume (Logical Operator)

- op_copy_services (Copy Services Operator)
- service (Service Operator)
- 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 **-1** 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 myds8000.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.

i5/0S

For the i5/OS, your personal profile is /home/<username>/dscli/security.dat.

OpenVMS system

For an OpenVMS operating system, the property value defaults to the logical name **SYS\$LOGIN**. As a result, your personal profile is [.dscli.profile]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. This command is not supported on DS6000 models.

After you create an authentication policy, use the manageauthpol command to configure it, and the **chauthpol** command to enable it.

```
► mkauthpol -type sas pol_name
```

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.

► mkuser -pw password -group group name [...] -pol pol name

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), Netware, or OpenVMS 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)
- 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)

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. This parameter is not supported on DS6000 models.

-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.

Example

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.

rmauthpol

The **rmauthpol** command allows you to remove an authentication policy. This command is not supported on DS6000 models.



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.

▶ — rmuser —			——User Name—	M
riiusei				
	└ -quiet_/	└─ -pol── pol_name─	_ "_"	

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.

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.

Using 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. 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.

▶ → setauthpol — -action —	setauthserver	►
	— addmap————	
	— rmmap———	
	— rmallmap———	
	— setmap———	
	— setsasuser——	
	— settruststore—	
	└─ setlocaladmin┘	

Each of the following sections shows the options for one of the listed **-action** parameters:

-action setauthserver



•	
e	xtuser—extusr1[,extusr2]dsgroup—dsgrp1[,dsgrp2]
►s	cope—user_resource_scope "_"
-actio	n rmallmap
►►— se	tauthpol— -action—rmallmapextgroup— <i>extgrp1[,extgrp2]</i>
►e	xtuser—extusr1[,extusr2]dsgroup—dsgrp1[,dsgrp2]
►s	cope—user_resource_scope pol_name
-actio	n setmap
►►— se	tauthpol— -action— setmap—extgroup— <i>extgrp1[,extgrp2]</i>
►e	xtuser—extusr1[,extusr2]dsgroup—dsgrp1[,dsgrp2]
►- <u>-</u> s	cope—user_resource_scope pol_name►◄
-actio	n setsasuser
► se	tauthpol— -action— setsasuser—username— namepw— password pol_name
-actio	n settruststore
► se	tauthpol— -action— settruststore polname ><
-actio	n setlocaladmin
► se	tauthpol— -action— setlocaladmin—username— <i>name</i> enabledisable
	l_name "

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.

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

Table 10. setauthpol action and parameter combinations

Table 10. setauthpol action and parameter combinations (continued)

Action	Parameter group 1	Parameter group 2	Effects
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", "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 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:

- 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 test2ibm my_policy1
```

The resulting output

Authentication policy my_policy1 successfully modified.

showauthpol

The **showauthpol** command displays detailed properties of a specified authentication policy. This command is not supported on DS6000 models.



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

DS_group	Ext_group
op_volume	grpa,grpb
op_copy_services	grpa,grpb

DS_group	Ext_user
admin	joe,bob
no_access	bob

Scope	Ext_group
Accounting	grpa
Purchasing	grpb

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

Ext_group	DS_group
grpb	op_volume,op_copy_services
grpa	op_volume,op_copy_services

Ext_user	DS_group
bob	admin,no_access
joe	admin

Ext_group	Scope
grpb	Purchasing
grpa	Accounting

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.

► showpass-

└ -pol name┘

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.

▶ — showuser —		— _— —User Name—	
SHOWUSCI		ober_name	
	└─ -pol─_ <i>pol_name</i> ─		

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.
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. This command is not supported on DS6000 models.

▶▶—	testauthpol— -username— user_name—	-pw— password-			
			└ -scope—	resource_scope_	
▶──	-group— dsgroup1[,dsgroup2,]	polname			

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 test2ibm 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.



Parameters

- S

(Optional) Displays the user names for the users who are currently logged on. The **-1** and **-s** parameters cannot be used together.

-1

(Optional) Displays details about the users who are currently logged on, including user name and authority groups that the user belongs to. The **-1** 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 will be 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)
- 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:

DSCIM

DS open application programming interface

DSCLI

DS Command Line Interface

DSGUI

DS Storage Manager Interface

HMTU

IBM Easy Tier Heat Map Transfer Utility

TPC IBM TotalStorage Productivity Center for Disk

TPC-R

IBM TotalStorage Productivity Center for Replication

Unknown

Unknown might be displayed for older versions of any of the client types 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	Туре
Testuser	service,op_copy_services	localPolicy	ABC_Co	DSGUI
Biguser	admin	localPolicy	*	DSCLI
Biguser	admin	localPolicy	*	DSCLI

Version	Address	Protocol
-	9.123.236.45	Internal
7.6.10.204	9.123.236.139	TLSv1.2
7.6.10.204	9.123.236.139	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

Indicates the client connection protocol, as follows:

None

The connection is not encrypted.

Internal

The connection is wholly contained within the HMC.

SSLv3

Indicates Secure Socket Layer version 3

TLSv1.0

Indicates Transport Layer Security version 1.0

TLSv1.1

Indicates Transport Layer Security version 1.1

TLSv1.2

Indicates 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. This command is not supported on DS6000 models.



Parameters

-s

(Optional) Provides the user name of the current user. The **-1** and **-s** parameters cannot be used together.

-1

(Optional) Provides details about the current user, including user name and authority groups that the user belongs to. The **-1** 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	*

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 disjoint 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.

1skeygrp

Displays a list of the key server encryption key group entries on the specified storage image.

1skeymgr

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. This command is not supported on DS6000 models.

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.

Iskeygrp

The **1skeygrp** command displays a list of the key server encryption key group entries on a specified storage image. This command is not supported on DS6000 models.



Parameters

-dev storage_image_ID

(Optional) Displays key group entries 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 encryption 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.

- S

(Optional). Displays only the attributes that are identified as short output. You cannot use the **-s** and the **-1** parameters together.

-1

(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) Specifies the state of the encryption group.

-reckeystate configured | newkeyveripend | newkeyauthpend | rekeyveripend | rekeyauthpend | recovauthpend | disabled | enableauthpend | disableauthpend | disableauthpend | Optional) Specifies encryption groups with the specified recovery key state.

-grpstatus critical | degraded | failed | normal | not_normal

(Optional) Displays the encryption key access status for all key servers that are associated with the specified encryption group. The value not_normal displays for all encryption groups whose encryption 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 encryption key group. The value not_normal displays for all encryption groups whose key server path summary status is not "normal" or the state is inactive.

Encryption_Group_ID ... | -

(Optional) Specifies the ID for the encryption 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

Invoking the lskeygrp command to display key server information.

dscli> lskeygrp

The resulting output

```
Date/Time: December 13, 2013 8:28:26 AM MST IBM DSCLI Version: 7.7.20.86 DS:
ID state reckeystate reckeydate datakeydate
1 accessible configured 09/28/2013 04/18/2013
```

Date/Time: December 13, 2013 8:28:26 AM MST IBM DSCLI Version: 7.7.20.86 DS: ID ... reckeydate datakeydate grpstatus keystatus label label2 1 ... 2013-09-28 2013-04-18 critical normal CompanyABC CompanyABC2

Report field definitions

ID Indicates the encryption group ID.

state

Indicates one of the following states of the encryption group:

accessible

The encryption group is accessible if it is configured and the storage image has the encryption key from the key server for the encryption group.

inaccessible

The encryption group is inaccessible if the storage image was unable to obtain the encryption key from the key server.

unconfigured

The encryption group is unconfigured if it has not been configured.

rekeying

The encryption group is accessible and rekeying if it is configured and the storage image has the encryption key from the key server for the encryption 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 encryption 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

The date of the last recovery key creation.

datakeydate

The date of the last data key creation. If the encryption group is unconfigured, then any displayed date is to be considered erroneous data.

grpstatus

Indicates one of the following values of the encryption key access status:

critical

Indicates that the encryption group has access to the encryption 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

Indicates that the encryption 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

Indicates that the encryption 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

Indicates that the encryption group has access to the encryption key on all key servers.

"_"

Indicates by the dash (-) that the encryption group state is unconfigured.

mgrstatus

Indicates one of the following values of the key server path access status:

critical

Indicates that at least one key server for this encryption 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

Indicates that at least two key servers for the specified encryption 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

Indicates that all key servers for this encryption group reported an access status of failed. Use the **showkeymgr** command with the –access parameter to determine the access status of each HMC.

normal

Indicates that all key servers for this encryption group reported an access status of normal.

"_"

Indicates by the dash (-) that the state of the encryption group is either unconfigured or all key servers report a state of inactive.

label


Indicates the label for the key server encryption 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 encryption 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

lskeymgr

The **lskeymgr** command displays a list of key servers that are on the storage complex and provides status information for those key servers. This command is not supported on DS6000 models.



Parameters

-s (Optional) Displays the key server IDs.

-1 (Optional) Displays the default output.

-serverport serverport_ID

(Optional) Displays the key servers that use the server port ID that you specify. The key server port ID is four or five decimal characters. For example, 8100 is a valid server port ID.

-state active | inactive

(Optional) Displays the key servers that are in the state that you specify.

-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 TKLM key server status:

dscli> lskeymgr

The resulting output

```
Date/Time:December 13, 2013 8:01:18 AM MST IBM DSCLI Version: 7.7.20.86 DS:IDStateStatusAddrServer Port1activenormaltklm.storage.tucson.ibm.com3801
```

Example

Invoking the lskeymgr command to display status for the key server whose status is not_normal:

dscli> lskeymgr -status not normal

The resulting output

```
        Date/Time:
        December
        13,
        2013
        8:01:18
        AM MST
        IBM
        DSCLI
        Version:
        7.7.20.86
        DS:

        ID
        State
        Status
        Addr
        Server
        Port

        2
        inactive
        -
        2001:0db8:85a3:08d3:1319:8a2e:0370:7334
        3801
```

Report field definitions

 \mathbf{ID}^{*}

The key server identification number. For example, 4.

State

Indicates one of the following states of the key server:

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 one of the following statuses of the key server path:

critical

Indicates that 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

Indicates that two or more HMCs have access to the specified key server, but at least one other HMC does not.

failed Indicates that neither HMC1 or HMC2 have access to the 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.

Use the ${\it showkeymgr}$ command with the –access parameter to determine the status of each HMC.

hmc2_degraded

Indicates that 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

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

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.

managekeygrp

The **managekeygrp** command allows you to manage an encryption key group.

The key servers use the key labels to locate a specific encryption key that is stored in the key server. With the introduction of the dual platform key server support, a second label was added 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, replacing any existing labels.

The DS8000 tests the connection to the encryption key that is stored on each key server once every eight hours. To test the encryption key access for a specific encryption 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) 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 encryption 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.

-action rekey | testaccess | updatecert

(Required) Specifies an action for the encryption key group. The valid values are as follows:

rekey

Specifies that a new key, as specified by the -key parameter, should be generated.

testaccess

Specifies that access to the encryption group key for the specified encryption group must be tested. Use the **showkeygrp –access** command to view the results.

updatecert

This option also specifies that a new key, as specified by the **-key** parameter, should be generated. It also updates the key client certificate from a Gen-1 to a NIST SP 800-131a compliant Gen-2 certificate that is used to authenticate with the key server.

Note: Before using the updatecert action, ensure that all key servers that are used by the DS8000 system contain the Gen-2 trust anchor certificate. Switching back to a Gen-1 certificate requires contacting IBM support.

-key data

(Optional) Use this parameter to specify that the encryption key group data key will be 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.

data specifies a random 256-bit encryption key that is generated by the key server and is used by the DS8000 system to wrap the encryption group key. This key is identified by one of the key labels that is specified by the -label or -label2 parameters.

-label key_label

(Optional) Use this parameter to specify the label for the encryption key group data key. This parameter is required when using the **-label2** parameter. You can enter a maximum of 64 ASCII characters for the label, which is used by a key server to identify a specific data key that the DS8000 system uses to wrap the encryption group key.

-label2 second_key_label

(Optional) This parameter is used by a second key server to identify a specific data key that the DS8000 system uses to wrap the encryption 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.

encryption_group_ID | -

(Required) Specifies the ID for the encryption group. The encryption group ID is a decimal number that ranges from 1 to N, where N is the maximum number of encryption 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.

Example

An invocation 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 encryption 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.

```
►→managekeymgr→ -action _____ setcert _____ key_server_ID _____ testaccess _____ -loc ___ location ____ "_" ____
```

Parameters

-action setcert testaccess

(Required) Specifies the location of the local certificate file.

setcert

Specifies the replacement of the existing anchor certificate that is used in secure communications with the key server.

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.

-loc location

(Optional) 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

Note: This parameter is required for the action values setcert.

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

An invocation 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 encryption 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 encryption 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.

Cance1

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 encryption 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 encryption 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.

encryption_group_ID

(Required) Specifies the encryption group ID for the encryption recovery key that you want to manage. The encryption group ID is a decimal number that ranges from 1 to N, where N is the maximum number of encryption 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 encryption group 1 has been verified, authorization pending.

mkkeygrp

The **mkkeygrp** command creates an entry for the key server encryption key group on the 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 encryption 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 key_label

(Required) Specifies the label for the encryption key group data key. You can enter a maximum of 64 ASCII characters for the label.

-label2 second_key_label

(Optional) Specifies the second label for the encryption key group data key. You can enter a maximum of 64 ASCII characters for the label.

Encryption_Group_ID

(Required) Specifies the ID for the encryption group. The encryption group ID is a decimal number that ranges from 1 to N, where N is the maximum number of encryption 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> mkkeygrp -dev IBM.2107-75FA120 -label MyCompany -label2 MyCompany2 1

The resulting output


The key server encryption 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 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 as appropriate.

To use an encryption 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 serverport_ID

(Optional) Specifies the key server port ID. The key server port ID is four or five 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.

-state active | inactive

(Optional) Specifies the state of the key manager, where

active Specifies that this key manager should be used to store encryption keys and that it should be checked periodically to verify the health of the key manager.

inactive

Specifies that this key manager should not be used to store encryption keys and that it should not be checked periodically.

-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_manager_ID

(Required) Specifies the key manager ID. The key manager 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 -addr 9.11.56.78 4


The resulting output

The key server 4 has been created successfully.

mkreckey

The **mkreckey** command is for users with security administrator authority. It is the first step in creating a new Encryption 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. This command is not supported on DS6000 models.



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 encryption 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.

encryption_group_ID

(Required) Specifies the encryption group ID for the new encryption recovery key that you want to

create. The encryption group ID is a decimal number that ranges from 1 to N, where N is the maximum number of encryption 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 encryption key group on a specified storage image. This command is not supported on DS6000 models.



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 encryption 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 removal confirmation prompt for this command.

encryption_group_ID

(Required) Specifies the ID for the encryption 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

An invocation example


dscli> rmkeygrp -dev IBM.2107-75FA120 1

The resulting output

Are you sure you want to delete the key server encryption group 1? y/n y The key server encryption group 1 has been deleted.

rmkeymgr

The **rmkeymgr** command removes a key server entry on the storage complex. This command is not supported on DS6000 models.



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 1

The resulting output

Are you sure you want to delete key server entry 1? y/n y The key server 1 has been deleted.

rmreckey

The **rmreckey** command allows you to remove an encryption recovery key. This command is not supported on DS6000 models.

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 encryption 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.

encryption_group_ID

(Required) Specifies the encryption group ID for the encryption recovery key that you want to deconfigure. The encryption group ID is a decimal number that ranges from 1 to N, where N is the maximum number of encryption 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> rmreckey -dev IBM.2107-75FA120 -1300861 1

The resulting output

The Recovery Key for encryption group 1 will be de-configured after authorization.

showkeygrp

The **showkeygrp** command displays detailed information for a specified encryption key group. This command is not supported on DS6000 models.



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 encryption 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 encryption group last assessed the encryption key on the key server.

Encryption_Group_ID

(Required) Specifies the ID for the encryption group.

Example

Invoking the showkeygrp command to display key server information

dscli> showkeygrp -dev IBM.2107-75FA120 1

The resulting output

```
Date/Time: December 13, 2013 9:01:20 AM MST IBM DSCLI Version: 7.7.30.86 DS:
IBM.2107-75YZ881
TD
              1
numranks
              1
numpools
              1
state
             configured
reckeystate configured
reckeydate 09/28/2013 19:10:38 MST
datakeydate 04/18/2013 16:27:35 MST
label
             CompanyABC
label2
              CompanyABC2
certificate GEN1
grpstatus
              critical
keystatus
              normal
```

Report field definitions

ID Indicates the encryption group ID.

numranks

Indicates the number of ranks that are configured in the specified encryption group.

numpools

Indicates the number of extent pools that are configured in the specified encryption group.

state Indicates one of the following states of the encryption group:

accessible

The encryption group is accessible if it is configured. The storage image has obtained the encryption key from the key server for the encryption group.

inaccessible

The encryption group is inaccessible if the storage image was unable to obtain the encryption key from the key server.

unconfigured

The encryption group is unconfigured if it has not been configured.

rekeying

The encryption group is accessible and rekeying if it has been configured. The storage image has obtained the encryption key from the key server for the encryption group and is in the process of rekeying.

reckeystate

Indicates one of the following states of the recovery key:

configured

A new recovery key has been 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 encryption group will be 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 encryption group is unconfigured, then any displayed date is to be considered erroneous data.

- **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. GEN1 is a legacy certificate that is not compliant with NIST SP 800-131a. GEN2 is a NIST SP 800-131a-compliant certificate that is introduced in Release 7.2.

grpstatus

Indicates one of the following access statuses of the encryption key:

critical

Indicates that the encryption 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 encryption group has access to the encryption key on two or more key servers, but not all key servers.

failed Indicates that the encryption group does not have access to the encryption key on any key server.

normal

Indicates that the encryption group has access to the encryption key on all key servers.

"-" Indicates by the dash (-) that the encryption group state is unconfigured.

mgrstatus

Indicates one of the following values of the key server path access status:

critical

Indicates that at least one key server for this encryption 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 encryption group reported an access status of normal or degraded.

failed Indicates that all key servers for this encryption group reported an access status of failed.

normal

Indicates that all key servers for this encryption group reported an access status of normal.

"-" Indicates by the dash (-) that the encryption group state is either unconfigured or all key servers report a state of inactive.

Example

Invoking the showkeygrp command to display encryption key access per key server (KeyMgr)

dscli> showkeygrp -access 1

The resulting output

 Date/Time:
 December
 11, 2013
 8:57:44
 AM
 MST
 IBM
 DSCLI
 Version:
 7.7.20.524
 DS:

 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 encryption group ID in ranges from 1 to the number of encryption 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 encryption 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:

уууу	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 encryption 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:

уууу	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 encryption key server has never failed to access the encryption key.

showkeymgr

The **showkeymgr** command displays detailed properties of a specified key server entry. You can also use this command to view when the hardware management console (HMC) last accessed the specified key server.

▶ — showkeymar — -		kev server ID	
▶ Showkeymgr			
L	– -access	L"_"	

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

The resulting output

Date/Time: December 11, 2013 8:57:44 AM MST IBM DSCLI Version: 7.7.20.524 DS: ID 1 State active Status hmc2_degraded addr tklm.storage.tucson.ibm.com port 3801 protocol none certificate -

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

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.

certificate

Indicates the certificate file. A dash (-) is displayed if no certificate was specified.

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		success	2013-12-11T06:37:46-700	2013-04-26T16:23:18-700
1		failure	2013-11-29T21:22:09-700	2013-12-11T06:37:47-700

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:

уууу	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:

уууу	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

This section contains commands that are used to manage network security on the DS8000 system using the DS CLI.


The following user access security commands are available:

- **chaccess** The **chaccess** command allows you to change one or more access settings of a hardware management console.
- **1saccess** The **1saccess** 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 via the Internet or a VPN connection. This command affects service access only and does not change access to the machine via the DS Command Line Interface.

-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 machine through the DS Storage Manager.

-modem enable | disable

(Optional) Specifies whether you want to enable or disable the modem dial-in and VPN initiation to the HMC.

-hmc 1 | 2 | all

(Optional) Specifies to which HMC you want the access settings to apply. hmc 1 specifies the primary HMC, and hmc 2specifies the secondary HMC. The default value allspecifies 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.

Isaccess

The **lsaccess** command displays the access settings and virtual private network (VPN) status 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 **-1** parameters together.

-1

(Optional) Displays default output plus additional attributes that are identified as long output. You cannot use the **-1** and the **-s** parameters together.

-hmc 1 | 2 | all

(Optional) A value of 1 specifies the primary HMC, and a value of 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

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 **lsaccess** command using the **-hmc** parameter.

Invoking the lsaccess command to display a list of the HMC access settings.

dscli> lsaccess -hmc all -l

The resulting output

hmc	cmdline	wui	modem	vpn
hmc1	enabled	enabled	enabled	disabled
hmc2	-	-	-	-

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 unavailable.

This field is for service access only and is not related to access to the machine 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 unavailable.

modem

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 unavailable.

vpn⁺

Indicates the connection status (enabled or disabled) of the outbound VPN connection. A status of enabled means a VPN connection exists, and a status of disabled means that a VPN connection does not exist.

Key:

- * Displayed when the **-s** parameter is specified.
- Displayed only when the -1 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 and for encrypting communications of applications that use VPN (Internet and modem) connections. 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 additional 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 IBM support personnel with special access rights can use the VPN connection.
- 3. The VPN connection is used when remote access is required by IBM Support personnel and no local onsite access to the machine is available.
- 4. The VPN connection can take from 2 to 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 issuing 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 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.

▶ → manageaccess → -ctrl → port1750 → -action	n enable	_
	$\Box_{disable}$ \Box_{-quiet}	

The following diagram shows the **manageaccess** command with the **-ctrl** parameter value set to **cim** | **gui** | **wui** | **ni**. The valid **-action** value for these controls is setsecurity. You must also specify the **-level** parameter.



Parameters

-ctrl cim | gui | wui | ni | vpn | port1750

cim This control specifies the inbound connection protocol to the CIM (SMI-S) server on the HMC or HMCs. Set this control to the 800131a level when all CIM clients that are accessing the CIM server on the HMC or HMCs 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.

Note: Setting this control causes the CIM server to reboot. All programs with CIM clients are forced to log on again.

gui This control specifies the inbound connection protocol to the DS Storage Manager GUI through browsers such as Internet Explorer or Firefox. Set this control to the 800131a level only when all browsers that are accessing the DS Storage Manager 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 DS Storage Manager GUI to reboot. All DS Storage Manager 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, TPC, TPC-R, 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 reboot. Any users who meet the specified security level requirements will automatically reconnect. However, all commands from the NI client programs, such as DS CLI, TPC, and TPC-R, will 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.

Note: Setting this control does not cause any server to reboot.

port1750

Specifies the inbound port 1750 to the DS8000 system from external legacy programs, such as the DS CLI, TPC, TPC-R, and others, that were created prior to Release 7.2. While 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 reboot.

-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 cim, 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: cim, 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 should be modified.

Notes:

- -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. If this is the case, 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.

showaccess

The **showaccess** command displays the access properties of a specified hardware management console (HMC).

	HMC ID		-
Silowaccess-			
	L "_"	l de la constante de	

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 this is the case, using hmc1 in this command still refers to the actual primary HMC.

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 **showaccess** command.

Invoking the showaccess command

dscli> showaccess hmc1

The resulting output

hmc	cmdline	wui	modem	vpn	port1750
hmc1	enabled	enabled	enabled	disabled	enabled

cim (security)	gui (security)	ni (security)	wui (security)
legacy	legacy	legacy	legacy

Report field definitions

Note: A dash (-) is displayed if the state (cmdline, wui, modem, vpn, port1750, cim [security], gui [security], ni [security], and wui [security]) is unknown. This result might occur if the client is not connected to the HMC.

hmc Indicates the HMC for which access properties are listed. hmc1 is the primary HMC, and hmc2 is the secondary HMC.

cmdline

Indicates the access setting (enabled or disabled) for the command-line shell on the HMC through Internet or VPN connection. This information is for service access only and is not related to access to the system through the DS command-line interface.

wui Indicates the access setting (enabled or disabled) for the HMC Web User Interface (WUI) on the HMC through Internet or VPN connection. This information is for service access only and has no effect on access to the system by the DS Storage Manager.

modem

Indicates the access setting (enabled or disabled) for the modem dial-in and VPN initiation to the HMC.

vpn Displays the connection status (enabled or disabled) of the outbound VPN connection.

port1750

Displays the status (enabled or disabled) of the legacy port 1750.

cim (security)

Indicates the security protocol that is allowed for the inbound connections to the CIM (SMI-S) 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.

gui (security)

Indicates the security protocol that is allowed for the inbound connections to the DS Storage Manager 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.

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. 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.

Internet Protocol Security (IPSec) commands

This topic contains commands that are used to maintain Internet Protocol Security (IPSec).

The DS8000 system supports IPSec connections that you can specify the connection by using a connection file. This connection file has the following format:

conn conn_ID

(Required) Specifies a connection definition with the name of *conn_ID*. This line is immediately followed by other parameters, one per line, without intervening blank lines. The connection definition ends by a blank line that follows all other parameters. The *conn_ID* is required by most IPSec commands except **mkipsec**, and is limited to the following characters: upper and lowercase alphabetic, numeric, dash (-), underscore (_), and period (.).

authby = pubkey | rsasig | psk | secret

pubkey | rsasig

Specifies the public key signature authentication, including **rsasig** (RSA digital signature). The default is **pubkey** if **authby**, **leftauth**, nor **rightauth** is specified.

psk | secret

Specifies the pre-shared key authentication.

Note: This parameter is deprecated, but still accepted, since two peers are not required to use the same authentication method. The **leftauth** and **rightauth** parameters should be used to specify the individual authentication methods.

auto = add | route | start

add

Specifies that the connection is loaded by the server, but disabled.

route

Specifies that the connection is loaded in such a way that if an attempt to connect is detected between the left and right peers, the connection is enabled.

start

Specifies that the connection is loaded by the server, and is enabled.

Note: This parameter is required in all connection definitions.

esp = cipher_suites

Specifies a list of ESP (Encapsulating Security Payload) encryption and integrity algorithm pairs used with a connection. The format of *cipher_suites* is a comma-separated list of encryption and integrity algorithms with the following format:

encryption-integrity[-dhgroup][-esnmode]

encryption

Specifies an encryption algorithm keyword.

integrity

Specifies an integrity algorithm keyword.

dhgroup

Specifies a Diffie-Hellman group keyword. If specified, a separate Diffie-Hellman exchange is used for CHILD_SA setup and rekeying.

esnmode

Specifies the extended sequence number support mode. The valid values are **esn** and **noesn**. The default is **noesn** if **esnmode** is not specified.

The following table lists the valid keywords for **encryption**, **integrity**, and **dhgroup**.

Table 11. Valid keywords for encryption, integrity, and dhgroup

Keyword	Description	IANA	IKE	ESP
3des	168 bit 3DES-EDE-CBC	3	x o g a	k
cast128	128 bit CAST-CBC	6	o g a	k
blowfish128 or blowfish	128 bit Blowfish-CBC	7	x o g a	k
blowfish192	192 bit Blowfish-CBC	7	x o g	k
blowfish256	256 bit Blowfish-CBC	7	x o g	k
null	Null encryption	11	-	k
aes128 or aes	128 bit AES-CBC	12	x o g a	k
aes192	192 bit AES-CBC	12	x o g a	k
aes256	256 bit AES-CBC	12	x o g a	k
aes128ctr	128 bit AES-COUNTER	13	x o g a	k
aes192ctr	192 bit AES-COUNTER	13	x o g a	k
aes256ctr	256 bit AES-COUNTER	13	x o g a	k
aes128ccm8 or aes128ccm64	128 bit AES-CCM with 64 bit ICV	14	x o g a	k
aes192ccm8 or aes192ccm64	192 bit AES-CCM with 64 bit ICV	14	x o g a	k
aes256ccm8 or aes256ccm64	256 bit AES-CCM with 64 bit ICV	14	x o g a	k
aes128ccm12 or aes128ccm96	128 bit AES-CCM with 96 bit ICV	15	x o g a	k
aes192ccm12 or aes192ccm96	192 bit AES-CCM with 96 bit ICV	15	x o g a	k
aes256ccm12 or aes256ccm96	256 bit AES-CCM with 96 bit ICV	15	x o g a	k
aes128ccm16 or aes128ccm128	128 bit AES-CCM with 128 bit ICV	16	x o g a	k
aes192ccm16 or aes192ccm128	192 bit AES-CCM with 128 bit ICV	16	x o g a	k
aes256ccm16 or aes256ccm128	256 bit AES-CCM with 128 bit ICV	16	x o g a	k
aes128gcm8 or aes128gcm64	128 bit AES-GCM with 64 bit ICV	18	x o g a	k
aes192gcm8 or aes192gcm64	192 bit AES-GCM with 64 bit ICV	18	x o g a	k
aes256gcm8 or aes256gcm64	256 bit AES-GCM with 64 bit ICV	18	x o g a	k

Keyword	Description	IANA	IKE	ESP
aes128gcm12 or aes128gcm96	128 bit AES-GCM with 96 bit ICV	19	x o g a	k
aes192gcm12 or aes192gcm96	192 bit AES-GCM with 96 bit ICV	19	xoga	k
aes256gcm12 or aes256gcm96	256 bit AES-GCM with 96 bit ICV	19	x o g a	k
aes128gcm16 or aes128gcm128	128 bit AES-GCM with 128 bit ICV	20	x o g a	k
aes192gcm16 or aes192gcm128	192 bit AES-GCM with 128 bit ICV	20	x o g a	k
aes256gcm16 or aes256gcm128	256 bit AES-GCM with 128 bit ICV	20	x o g a	k
aes128gmac	Null encryption with 128 bit AES-GMAC	21	-	k
Aes192gmac	Null encryption with 192 bit AES-GMAC	21	-	k
aes256gmac	Null encryption with 256 bit AES-GMAC	21	-	k
camellia128 or camellia	128 bit Camellia-CBC	23	o g a	k
camellia192	192 bit Camellia-CBC	23	o g a	k
camellia256	256 bit Camellia-CBC	23	o g a	k
camellia128ctr	128 bit Camellia-COUNTER	24	o g a	k
camellia192ctr	192 bit Camellia-COUNTER	24	o g a	k
camellia256ctr	256 bit Camellia-COUNTER	24	o g a	k
camellia128ccm8 or camellia128ccm64	128 bit Camellia-CCM with 64 bit ICV	25	o g a	
camellia192ccm8 or camellia192ccm64	192 bit Camellia-CCM with 64 bit ICV	25	o g a	
camellia256ccm8 or camellia256ccm64	256 bit Camellia-CCM with 64 bit ICV	25	o g a	
camellia128ccm12 or camellia128ccm96	128 bit Camellia-CCM with 96 bit ICV	26	o g a	
camellia192ccm12 or camellia192ccm96	192 bit Camellia-CCM with 96 bit ICV	26	o g a	
camellia256ccm12 or camellia256ccm96	256 bit Camellia-CCM with 96 bit ICV	26	o g a	
camellia128ccm16 or camellia128ccm128	128 bit Camellia-CCM with 128 bit ICV	27	o g a	
camellia192ccm16 or camellia192ccm128	192 bit Camellia-CCM with 128 bit ICV	27	o g a	
camellia256ccm16 or camellia256ccm128	256 bit Camellia-CCM with 128 bit ICV	27	o g a	

 Table 11. Valid keywords for encryption, integrity, and dhgroup (continued)

Key:

IANA IANA (Internet Assigned Numbers Authority) encryption number

- x Default strongSwan built-in cryptographic library
- OpenSSL (Open SSL project) cryptographic library
- g gcrypt (GNU cryptographic) cryptographic library
- a AF_ALG user-space cryptographic API for the Linux 2.6.38 kernel or newer
- k Linux 2.6 kernel

Keyword	Description	IANA	IKE	ESP	Info
md5	MD5 HMAC	1	96 bit	96 bit	
md5_128	MD5_128 HMAC	6	n/a	128 bit	x
sha1 or sha	SHA1 HMAC	2	96 bit	96 bit	
sha1_160	SHA1_160 HMAC	7	n/a	160 bit	x
aesxcbc	AES XCBC	5	96 bit	96 bit	
sha2_256 or sha256	SHA2_256_128 HMAC	12	128 bit	128 bit	t
sha2_384 or sha384	SHA2_384_192 HMAC	13	192 bit	192 bit	
sha2_512 or sha512	SHA2_512_256 HMAC	14	256 bit	256 bit	
sha2_256_96 or sha256_96	SHA2_256_96 HMAC	р	96 bit	96 bit	t

Table 12. Integrity algorithm keywords for esp or ike

Key:

IANA IANA (Internet Assigned Numbers Authority) integrity number

- x Requires Linux 2.6.33 kernel or newer
- t Before Linux 2.6.33, the Linux kernel incorrectly used 96 bit truncation for SHA-256
- **p** strongSwan uses the value of 1026 from the IANA private use range

Table 13. Diffie-Hellman group keywords for esp or ike

Keyword	DH Group	Modulus	IKE
modp756	1	768 bits	m o g
modp1024	2	1024 bits	m o g
modp1536	5	1536 bits	m o g
modp2048	14	2048 bits	m o g
modp3072	15	3072 bits	m o g
modp4096	16	4096 bits	m o g
modp6144	17	6144 bits	m o g
modp8192	18	8192 bits	m o g

Table 14. Modulo Prime Groups with Prime Order Subgroup

Keyword	DH Group	Modulus	Subgroup	IKE
modp1024s160	22	1024 bits	160 bits	m o g
modp2048s224	23	2048 bits	224 bits	m o g
modp2048s256	24	2048 bits	256 bits	m o g

Table 15. Elliptic Curve Groups

Keyword	DH Group	Modulus	IKE
ecp192	25	192 bits	0
ecp224	26	224 bits	0
ecp256	19	256 bits	0
ecp384	20	384 bits	0
ecp521	21	521 bits	0

Key:

- m GMP (GNU Multi-Precision) big number library
- o OpenSSL (Open SSL project) cryptographic library
- g gcrypt (GNU cryptographic) cryptographic library

Note: The complete list of IANA transform type numbers can be found at Internet Assigned Numbers Authority website (www.iana.org/)

ike = cipher_suites

Specifies a list of IKE/ISAKMP (Internet Key Exchange/Internet Security Association and Key Management Protocol) encryption, integrity, and Diffie-Hellman algorithms that are used with a connection. The format of *cipher_suites* has the following format:

encryption-integrity-dhgroup

encryption

Specifies an encryption algorithm keyword.

integrity

Specifies an integrity algorithm keyword.

dhgroup

Specifies a Diffie-Hellman group keyword.

Note: The keywords for encryption, integrity, and dhgroup are listed in the tables under esp.

keyexchange = ike | ikev2

ike

Specifies the protocol to be used to initialize a connection. The default is **ike** if **keyexchange** is not specified, and is equivalent to **ikev2**.

ikev2

Specifies that the IKE version 2 protocol is to be used to initialize the connection.

Note: The IKE version 1 protocol (ikev1) is not supported for customer-specified connections.

type = tunnel | transport

tunnel

Specifies a host-to-host, host-to-subnet, or subnet-to-subnet IPSec tunnel mode. The default is **tunnel**, if type is not specified.

transport

Specifies a host-to-host IPSec transport mode.

The following keywords are defined in terms of connection left and right endpoints or peers. The left connection endpoint is considered to be the local peer endpoint that is associated with the DS8000 HMC, and the following documentation implies this assumption. The right connection endpoint is considered to be the remote peer endpoint.

left/right = *ip_address* | *fqdn* | %*any* | %*defaultroute*

ip_address

Specifies the peer's IP address in either IPv4 or IPv6 format.

fqdn

Specifies the peer's IP address as a Fully Qualified Domain Name.

%any

When used with the **right** keyword, specifies the remote peer's IP address that might be any IP address.

%defaultroute

When used with the **left** keyword, specifies the local peer's IP address.

Note: The default value for the **left** keyword is **%defaultroute** and the default for the **right** keyword is **%any**.

leftauth/rightauth = pubkey | psk

pubkey

Specifies public key signature authentication that includes RSA digital signature or Elliptic Curve DSA signature. The default is **pubkey** if **authby**, **leftauth**, nor **rightauth** is specified.

psk

Specifies pre-shared key authentication.

leftcert/rightcert = cert_name

cert_name

Specifies the peer's x.509 certificate's file name. The certificate file must be in PEM or DER format, and must be imported with the **mkipseccert** command.

The DS8000 HMC IPSec function is provided by strongSwan version 4.6.1. While the preceding connection definition keywords are supported by the DS CLI, other possible keywords might be required for your specific environment. See strongSwan website (www.strongswan.org) for more information.

The following Internet Protocol Security (IPSec) commands are available:

chipsec

The chipsec command modifies an existing Internet Protocol Security (IPSec) connection.

lsipsec

The **lsipsec** command displays a list of defined Internet Protocol Security (IPSec) connection configurations.

lsipseccert

The lsipseccert command displays a list of Internet Protocol Security (IPSec) certificates.

mkipsec

The **mkipsec** command creates an Internet Protocol Security (IPSec) connection by importing an

Internet Protocol Security (IPSec) connection configuration file that contains a connection definition to the hardware management console (HMC).

mkipseccert

The **mkipseccert** command imports an Internet Protocol Security (IPSec) certificate to the DS8000 system.

rmipsec

The **rmipsec** command deletes an Internet Protocol Security (IPSec) connection from the IPSec server.

rmipseccert


The **rmipseccert** command deletes an Internet Protocol Security (IPSec) certificate from the hardware management console (HMC).

setipsec

The setipsec command manages the Internet Protocol Security (IPSec) connections.

chipsec

The **chipsec** command modifies an existing Internet Protocol Security (IPSec) connection. This command is not supported on DS6000 models.



Parameters

-enable | -disable

(Optional) Specifies whether you want to enable or disable the IPSec connection. When you specify the **-enable** parameter, an attempt is made to establish the IPSec connection. When you specify the **-disable** parameter, the IPSec connection ends.

Notes:

- 1. The -enable and -disable parameters cannot be used together.
- 2. The value of auto in the connection configuration file also influences whether the connection is enabled or disabled. For example, if the connection configuration file is defined with auto=start, the IPSec server attempts to establish the connection when the connection is created, or whenever the IPSec server is restarted or updated.
- **3**. In an active connection, both peers must be sending and receiving data. An enabled connection might not be active if it is waiting for the peer to activate the other end of the connection.

-hmc 1 | 2 | all

(Optional) Specifies the HMC that you want to import the IPSec connection configuration settings:

- -hmc 1 specifies the primary HMC
- -hmc 2 specifies the secondary HMC
- all is the default value and it 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.

conn_ID ... | -

(Required) Specifies the IPSec connection IDs that you want the connection configuration settings to apply to. 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 chipsec command:

dscli> chipsec -enable connection1


The resulting output:

IPSec connection connection1 on hmc 1 successfully enabled.

lsipsec

The lsipsec command displays a list of defined Internet Protocol Security (IPSec) connections. This command is not supported on DS6000 models.





Parameters

- -s (Optional) Displays the attributes that are identified as short output.
- -1 (Optional) Displays the default output and attributes that are identified as long output.

-hmc 1 | 2 | all

(Optional) Specifies the HMC that you want to view the connection status for. "-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.

-state enabled | disabled

(Optional) Specifies whether you want to limit the output to only one of the two connection states (enabled or disabled).

conn ID ... | -

(Optional) Specifies an IPSec connection ID for which you want information to be displayed. 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 **lsipsec** command using the **-hmc** parameter.

Invoking the lsipsec command

dscli> lsipsec -hmc 1

The resulting output

ID	hmc	State
alpha	hmc1	Enabled
beta	hmc1	Disabled
gamma	hmc1	Enabled

Report field definitions

- **ID** Displays the identifiers and connection names of existing IPSec connections on the specified HMC.
- hmc Displays the hardware management console (HMC) identifier.
- **State** Indicates the state of the identified IPSec connection on the specified HMC. One of the following states is displayed:

Connecting

The IPSec server is establishing a connection.

Creating

The IPSec server is creating entries that are required to establish a connection.

Deleting

The IPSec server is deleting the connection.

Disabled

The IPSec connection is disabled.

Enabled

The IPSec connection is enabled.

Rekeying

The IPSec server is updating the encryption keys of the connection.

Routed

The IPSec connection is defined, but is not enabled until communications are detected by either end of the connection.

Server Down

The IPSec server is not running. This implies that all connections are disabled.


Unknown

Either the connection to the IPSec server is not known or the connection query response from the IPSec server was not understood by the DS CLI. A connection might be in the *Unknown* state if the values in the connection definition are not understood by the IPSec server. If a connection definition is not understood by the IPSec server, any subsequently created connections will also be in the *Unknown* state until the connection definition that is causing the original *Unknown* state is corrected.

"-" The connection is known to the IPSec server, but the state is not understood by the DS CLI.

Isipseccert

The **lsipseccert** command displays a list of Internet Protocol Security (IPSec) certificates. This command is not supported on DS6000 models.



Parameters

- -s (Optional) Displays the attributes that are identified as short output.
- -1 (Optional) Displays the default output and attributes that are identified as long output.

-hmc 1 | 2 | all

(Optional) Specifies the HMC that you want to view the certificate status for. "-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.

cert_ID ... | -

(Optional) Specifies the IPSec certificate ID that you want information to be displayed for. 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 **lsipseccert** command.

Invoking the lsipseccert command

dscli> lsipseccert

The resulting output

ID	hmc	HasKey
CertA.pem	hmc1	Yes
CertA.pem	hmc2	Yes
MyCert.pem	hmc1	No
MyCert.pem	hmc2	No

Report field definitions

ID Displays the identifier of an existing IPSec connection on the specified HMC.


hmc Displays the hardware management console (HMC) identifier.

HasKey

Indicates whether the certificate has an associated private key.

mkipsec

The **mkipsec** command creates an Internet Protocol Security (IPSec) connection on the DS8000 system by importing a configuration file that contains a connection definition to the hardware management console (HMC). IPSec connection file formats are described in the Internet Protocol Security (IPSec) tasks. This command is not supported on DS6000 models.



Parameters

```
-conn conn_path
```

(Required) Specifies the local directory path for the IPSec connection file.

-hmc 1 | 2 | all

(Optional) Specifies the HMC that you want to import the IPSec connection file to. "-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.

-phrase pass_phrase

(Optional) Specifies a pass phrase that is used to create an IPSec connection. Either the **-phrase** or the **-phrasefile** parameter must be used if the "authby", "leftauth", or "rightauth" values are set to "psk" or "secret".

Note: The -phrasefile and -phrase parameters cannot be used together.

-phrasefile phrase_path

(Optional) Specifies the path to a file that contains the pass phrase that is used to create an IPSec connection. Either the **-phrase** or the **-phrasefile** parameter must be used if the "authby", "leftauth", or "rightauth" values are set to "psk" or "secret".

Note: The -phrasefile and -phrase parameters cannot be used together.

-force

(Optional) Specifies that the IPSec connection in the specified connection file replaces any existing connection that has the same name.

Example

Invoking the mkipsec command

dscli> mkipsec -conn C:\temp\my_connections.conn

The resulting output


IPSec connection connection1 on hmc1 successfully created.

Contents sample of connection file "my_connections.conn"

```
conn connection1
authby=psk
auto=start
left=9.12.133.155
right=9.12.212.17
type=tunne1
keyexchange=ikev2
esp=aes128-sha256
```

mkipseccert

The **mkipseccert** command imports an Internet Protocol Security (IPSec) certificate to the DS8000. This command is not supported on DS6000 models.



Parameters

-cert cert_path

(Required) Specifies the local directory path for the IPSec certificate file. The certificate ID that is used in other IPSec certificate commands is the file name of the file in the specified certificate file path. The certificate name, like a connection name, is restricted to upper and lowercase alphabetic, numeric, dash (-), underscore (_), and period (.) characters. The certificate file must be in PEM or DER format.

-hmc 1 | 2 | all

(Optional) Specifies the HMC that you want to import certificate files to. "-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.

-key key_path

(Optional) Specifies the local directory path for the private key file. The private key file must be in PEM or DER format.

-phrase pass_phrase

(Optional) Specifies a pass phrase that might be required with a private key. This parameter is only valid when used with the **-key** parameter, and it is optional.

Note: The -phrasefile and -phrase parameters cannot be used together.

-phrasefile phrase_path

(Optional) Specifies the path to a file that contains the pass phrase that might be required with a private key. This parameter is only valid when used with the **-key** parameter, and it is optional.

Note: The -phrasefile and -phrase parameters cannot be used together.

-force

(Optional) Specifies that the specified IPSec certificate file overwrites any existing certificate with the same file name.

Example

Invoking the mkipseccert command


dscli> mkipseccert -cert C:\temp\CertA.pem

The resulting output

IPSec certificate CertA.pem on hmc1 successfully added. IPSec certificate CertA.pem on hmc2 successfully added.

rmipsec

The **rmipsec** command deletes an Internet Protocol Security (IPSec) connection definition from the IPSec server. This command is not supported on DS6000 models.



Parameters

-hmc 1 | 2 | all

(Optional) Specifies the HMC that you want to delete the connection definition 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.

-quiet

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

conn_ID ... | -

(Required) Specifies the IPSec connection IDs 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 rmipsec command


dscli> rmipsec -quiet connection1

The resulting output

IPSec connection connection1 on hmc1 successfully deleted.

rmipseccert

The **rmipseccert** command deletes an Internet Protocol Security (IPSec) certificate from the hardware management console (HMC). This command is not supported on DS6000 models.



Parameters

-hmc 1 | 2 | all

(Optional) Specifies the HMC that you want to delete the certificate 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.

-quiet

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

```
cert_ID ... | -
```

(Required) Specifies the IPSec certificate ID 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 rmipseccert command


dscli> rmipseccert connection1.pem

The resulting output

```
IPSec certificate connection1.pem on HMC1 successfully removed. IPSec certificate connection1.pem on HMC2 successfully removed.
```

setipsec

The **setipsec** command allows you to manage Internet Protocol Security (IPSec) controls. This command is not supported on DS6000 models.



Parameters

-action enable | disable

(Required) Specifies the action that you want to apply to the specified IPSec control. When you specify **-action** *enable* **-ctrl** *server*, an attempt is made to start the IPSec server. At least one connection configuration must first be imported using the **mkipsec** command before the server will start. When you specify **-action** *disable* **-ctrl** *server*, the IPSec server is stopped. All connections are terminated when the server is disabled.

-ctrl server

(Required) Specifies the IPSec control that you want the specified actions to apply to.

-hmc 1 | 2 | all

(Optional) Specifies the HMC that you want to apply the IPSec control action to. "-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.

-quiet

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

Example

Invoking the setipsec command

dscli> setipsec -action enable -ctrl server

The resulting output

The enable action for IPSec server on hmc1 has been performed.

Application key commands

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

The following application key commands are available:

app1ykey

Activates the licenses on your storage system.

1skey 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. For additional information about downloading the XML file and activating the LMC keys, see "Activating your machine and feature licenses using the DS CLI" on page 30.

```
►►— applykey-
```

```
y_____storage_image_ID_____storage_image_ID_____
```

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 unit.

The **1skey** command only displays the keys that are installed. Refer to the *IBM System Storage DS8000 Introduction and Planning Guide* for more information on how to choose which keys are needed and how to acquire them.

```
► lskey storage_image_ID
```

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), (not available for DS6000 models)	25	All	
HyperPAV (not available for DS6000 models)	On	CKD	
Metro/Global Mirror (MGM),	25	All	
Metro Mirror (MM), (not available for DS6000 models)	25	All	
Remote mirror and copy (RMC)	25	All	
Remote mirror for z/OS (RMZ)	25.1	CKD	
DS8000 I/O Priority Manager	25	All	
Operating Environment (OEL)	45	All	
IBM Easy Tier	On	All	
IBM FlashCopy SE (not available for DS6000 models)	25	All	
High Performance FICON for System z (zHPF) (not available for DS6000 models)	On	CKD	
RMZ Resync (not available for DS6000 models)	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.

Physical resource information commands

This section contains commands that are used to view information about the physical resources in your storage complex.

The following physical resource information commands are available:

- **1sda** Displays a report that lists the device adapters (DA) for each storage image and the status of each device adapter in the list. This command is not supported on DS6000 models.
- **1sddm** 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. This command is not supported on DS6000 models.

1shba Displays a report that lists the storage image host adapters and status information for each host adapter in the list. This command is not supported on DS6000 models.

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 **1sda** 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. This command is not supported on DS6000 models.



Parameters

- S

(Optional) Displays only device adapter IDs. You cannot use the -1 and the -s parameters together.

-1

(Optional) Displays the default output, plus the I/O enclosure and device adapter locations, feature codes, and interface IDs. You cannot use the **-1** and the **-s** parameters together.

-encl enclosure_ID

(Optional) Displays the device adapters that are associated with the specified processor complex 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 -1 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

FC	Server	DA pair	Interfs	DA_Pair_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

 \mathbf{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.ppsssss*.

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, 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.

WWN+

Indicates the World Wide 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 **-1** parameter is specified.

Isframe

The **lsframe** command displays a list of frame enclosures for a storage image. This command is not supported on DS6000 models.



Parameters

- -s Displays the rack enclosure ID. You cannot use the -1 and the -s parameters together.
- -1 Displays default output plus the frame ID and location of the frame enclosure. You cannot use the -1 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 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 **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

 \mathbf{ID}^{*}

Identifies the unique identifier of the frame enclosure.

frm#

Identifies the frame number within a storage unit that contains the specified frame enclosure.

Frame⁺

Identifies the unique identifier of the storage unit equipment frame that contains the specified frame enclosure.

Loc⁺

Identifies the frame enclosure location.

The location format is *Uttt.mmm.ppsssss*.

Key:

- * Displayed when the **-s** parameter is specified.
- + Displayed only when the **-1** parameter is specified.

Isstgencl

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 -1 and the -s parameters together.
- -1 (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 -1 and the -s parameters together.

-state normal | not_normal

(Optional) Displays all the storage enclosures that are associated with the specified storage unit 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 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

\mathbf{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 unit that contains this storage enclosure.

A " - " value is displayed for a DS6000 model.

Enclnum⁺

Indicates the identity of a storage enclosure within a storage unit frame.

Loc⁺

Indicates the storage enclosure location by identifying the storage unit frame that contains the storage enclosure. The location format is *Utttt.mmm.ppsssss*.

\mathbf{FC}^{\dagger}

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 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 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 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.

Note: For a DS6000 model, this condition changes to *inter failed* if the storage enclosure is found to be in good condition again.

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.

Key:

- * Displayed when the **-s** parameter is specified.
- + Displayed only when the **-1** parameter is specified.

lshba

The **1shba** 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 -1 and the -s parameters together.
- -1 (Optional) Displays the default output plus the host adapter feature code and interface IDs. You cannot use the -1 and the -s parameters together.

-encl enclosure_ID

(Optional) Specifies that the system displays host adapters that are associated with a common processor complex 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.ppsssss.

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.

lsddm

The **1sddm** 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 -1 parameters together.
- -1 (Optional) Displays the default output. You cannot use the -s and -1 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 displays DDMs that are associated with a specified state.

-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 | flash | nl | sata | ssd

(Optional) Displays the DDMs that are associated with the specified disk class.

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

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 **1sddm** command by using the **-1** parameter.

Invoking the lsddm command

dscli> lsddm -1 IBM.2107-75FA120

The resulting output

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	encrypt
1000	FCAL	2	array member	S123	1	normal	ENT	supported
1000	FCAL	2	array member	S123	2	normal	SATA	supported

Report field definitions

 \mathbf{ID}^{*}

Indicates the system-assigned unique identifier for the DDM.

Mode1

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).

Firmwarelevel

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 that the specified DDM resides on. 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. One of the following values are displayed:

- FC-AL
- 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
- Unconfigured
- Spare required
- · Spare not required
- Array member

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 has been 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
- Flash = Indicates high-performance flash devices.
- NL = Indicates near-line and represents ATA (FATA) disk drives
- SATA = Indicates high capacity SATA disk drives.
- SSD = Indicates solid-state devices.

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.

Key:

- * Displayed when the **-s** parameter is specified.
- + Displayed only when the **-1** parameter is specified.

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.
- **1svpn** Displays a report that lists the outbound VPN information such as management consoles and connection status. This command is not supported on DS6000 models.
- **setvpn** Used when remote access is required by IBM Support personnel and there is no local on-site access to the machine.
- 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 only be created or deleted by service personnel.
- The SNMP settings only apply to 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) This parameter is not currently supported.

- -emailaddr new_email_list
 (Optional) This parameter is not currently supported.
- -emailrelay on | off

(Optional) This parameter is not currently supported.

- -emailrelayaddr new_relay_ip_list
 (Optional) This parameter is not currently supported.
- -emailrelayhost new_relay_host_name
 (Optional) This parameter is not currently supported.
- -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 whenloginprompt is set to off.

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). This command is not supported on DS6000 models.


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 IBM support personnel with special access rights can use the VPN connection.
- 2. The **setvpn** command is used when remote access is required by IBM 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

Secure connection started.

lsvpn

The **lsvpn** command displays a report that lists the outbound VPN information such as management consoles and connection status. This command is not supported on DS6000 models.



Parameters

- S

(Optional) Displays the management console (SMC) and connection status. You cannot use the **-1** and **-s** parameters together.

-1

(Optional) Displays the management console (SMC) and connection status. You cannot use the **-1** 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 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 specified to start your session is the one that is listed in the report as being 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 identified by *smc1*. The console that you have specified 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 that is listed in the report as being 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.

▶ → — showsp-

Parameters

There are no parameters for this command.

Notes:

- 1. The SNMP settings only apply to 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. <i>xxx</i> .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

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 and disabled 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 " - " 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

The number of key servers that are supported.

loginprompt

Indicates the following values of the login message prompt:

0n

Indicates that the login message prompt is enabled.

0ff

Indicates that the login message prompt is disabled.

"_"

Indicates that the message prompt is not supported.

promptmessage

Indicates the login message prompt. If "-" displays, the promptmessage option is not supported.

Storage unit configuration commands

This section contains commands that are used to configure a storage unit.

A storage unit is a single, physical storage subsystem. A storage complex is a configuration of one or more storage units that is managed by a management console. If you have one DS8000 machine, then you have a single storage complex containing a single storage unit.

A storage image is a partitioning of a storage unit that provides emulation of a virtual storage server. You can configure more than one storage image on a storage unit.


The following storage unit configuration commands are available:

- **chsu** Changes the description and name you have associated with a specified storage unit. For DS8000, you can also change the remote power control mode on the storage unit.
- **Issu** Generates reports that allow you to view details about your storage units.

showsu Generates reports that allow you to view details about your storage units.

chsu

The **chsu** command modifies a storage unit. You can also use this command to power on and power off a DS8000 storage unit.



Parameters

-pwrmode manual | auto | zseries

(Optional) Sets a requested remote power control mode on the storage unit.

manual

Indicates that the storage facility power-on and power-off sequences are performed based on the manual power on and off controls.

auto

A storage facility power-on sequence is performed when external power first becomes available to the first rack of a storage facility (for example, when standby power is first 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 unit might be queued. This queuing can result in a loss of response on other functions that access the storage unit when accessed by the CLI.

-pwron

(Optional) Turns on power to the storage unit. 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 unit. 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 unit. 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 unit. 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 time period.

disable

Energy Report readings are averaged over a 5-minute time period

Note: IBM recommends disabling the Energy Report test mode at test completion.

storage_unit_ID | -

(Required) Accepts the fully qualified storage unit ID. The storage unit 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 unit IBM.2107-75FA120 successfully modified.

Issu

The **1ssu** command displays a list of storage units in a storage complex. You can use this command to look at the status and other properties of each storage unit in the list.

-s- -power on storage_unit_ID _____

Parameters

- -s (Optional) Displays only the storage unit ID. You cannot use the -1 and the -s parameters together.
- -1 (Optional) Displays default output plus the power mode and storage unit description. You cannot use the -1 and the -s parameters together.

-power on | off

(Optional) Displays only the storage units in the specified power state.

```
storage unit ID ... | -
```

(Optional) Displays storage units with the specified storage unit IDs. A storage unit 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 using 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	921	3007ACF 3012399 E0	On	Local	Test
SU 2	IBM.2107- 75FA120	921	3007ACF 3045699 E0	On	Local	Produc- tion
SU 3	IBM.2107- 75FA130	921	3007ACF 3078999 E0	On	Local	Backup

Report field definitions

Name

Indicates the user-defined name for each storage unit found within the storage complex. This value is " - " if you have not assigned a name to a storage unit.

ID Indicates the storage unit ID which 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.

Mode1

Indicates the model number of the storage unit.

WWNN

Indicates the World Wide Node Name for the listed storage unit. This value is " - " if the WWNN is not known

Pw State

Indicates the power status of the listed storage unit. One of the following values is displayed:

On Indicates the storage unit power is on.

0ff

Indicates the storage unit power is off.

Turning On

Indicates the storage unit power is in the process of turning on.

Turning Off

Indicates the storage unit power is in the process of turning off.

Power Exception

Indicates that storage unit power is on, but online operation is not possible due to a power fault in one of the storage unit frames.

Pw Mode

Indicates the power control mode in effect for the listed storage unit. One of the following values is displayed:

Loca1

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 unit. This value is displayed as a " - " if no description has been assigned.

showsu

The showsu command displays detailed properties of an individual storage unit.



Parameters

storage_unit_ID

(Required) Specifies the storage unit ID. A storage unit 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 Stor- age Unit	This is my DS Stor- age 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 unit. This value is " - " if you have not assigned a name to a storage unit.

Desc

Indicates the description that you assigned for the designated storage unit. This value is displayed as a " - " if no description has been assigned.

ID Indicates the storage unit ID which consists of the manufacture name, machine type, and serial number.

Mode1

Identifies the model number of the designated storage unit.

WWNN

Indicates the World Wide Node Name for the listed storage unit. This value is " - " if the WWNN is not known.

Config

Indicates the internal I/O interface configuration for the storage unit. 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 unit 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 unit. One of the following values is displayed:

On Indicates the storage unit power is on.

0ff

Indicates the storage unit power is off.

Turning On

Indicates the storage unit power is in the process of turning on.

Turning Off

Indicates the storage unit power is in the process of turning off.

Power Exception

Indicates that storage unit power is on, but online operation is not possible due to a power fault in one of the storage unit frames.

Pw Mode

Indicates the power control mode in effect for the listed storage unit. 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 configured on the storage unit.

MTS

(DS8000 models) Indicates the order type of the storage unit. The order type and the machine type of the storage unit is the same on all storage units ordered before release 2.4. After release 2.4, the order type varies according to the warranty periods associated with the storage unit.

(DS6000 models) No value is specified for a DS6000 model.

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, measured over a 5 minute time 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, measured over a 5 minute time 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, measured over a 5 minute time 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, measured over a 5 minute time 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.

Storage image configuration commands

This section contains commands that are used to configure 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 for the storage image, changes the description and name that you have 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.
- **1ssi** Displays a list of storage images in a storage complex. These commands requires that you use the storage image WWNN, which is displayed for each storage image when you use the **1ssi** 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.



Parameters

-essnetcs y | n

(Optional) Enables or disables the storage complex ESSNet user interface to invoke Copy Services operations for the storage image. y (yes) is the default.

-volgrp volume_group_ID

(Optional) Specifies the ESSNet Copy Services type volume group that contains the logical volumes that are eligible for control by Copy Services operations, when the **-essnetcs y** parameter is used. All logical volumes are eligible for control by Copy Services operations if the **-essnetcs y** parameter and the volume group ID are not specified.

The **-volgrp** 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*.

```
-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 3 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**** machines 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 only controls 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, provided that the IBM Easy Tier LIC feature is active.

automode

Easy Tier monitors only those volumes 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.

Note: Only volumes that are monitored by Easy Tier can participate in Easy Tier cooperative caching.

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 on this storage facility image.

enable

Specifies that the Easy Tier Heat Map Transfer feature is enabled.

disable

Specifies that the Easy Tier Heat Map Transfer feature is not enabled.

-iopmmode disable | monitor | monitorsnmp | manage | managesnmp

(Optional) Specifies the I/O priority management mode, which could be one of the following values:

disable

Specifies that the I/O priority manager function is disabled.

monitor

Specifies that resources associated with performance groups that specify I/O management are monitored, but not managed.

monitorsnmp

Specifies the same condition as *monitor*, but with added SNMP trap support.

manage

Specifies that resources associated with performance groups that specify I/O management are managed.

managesnmp

Specifies the same condition as *manage*, but with added SNMP trap support.

Note: The default I/O priority management mode is managed mode.

storage_image_ID | -

(Required) Specifies the storage image ID, which consists of the values for 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.

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 IBM Service.

- 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, DS8000 Storage Management GUI, and CIM 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

Diagnostic control oddlite is successfully submitted.

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 Diagnostic control saveuilogs is successfully submitted. It will take some time for the system to complete the request.

Isserver

The **1sserver** 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.

► lsserver Server_ID ... J

Parameters

-s (Optional) Displays only the server ID. You cannot use the -1 and the -s parameters together.

-1 (Optional) Displays the default output and the state of the servers. You cannot use the -1 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 using the **-1** 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

 \mathbf{ID}^{\star}

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. For a DS6000, this field always reports a dash.

Image Name

Indicates the image name for the designated storage server. For a DS6000, this field always reports a dash.

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 -1 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.

— storage_image_ID—		
storuge_image_ib		
	• • •	-

Parameters

- S

(Optional) Displays only the storage image IDs. You cannot use the **-1** and the **-s** parameters together.

-1

(Optional) Displays the default output, ESSNet, volume group, and storage image description. You cannot use the **-1** 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 unit. A storage unit 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 **-1** parameter. There is a difference in the input values when you use the **su_ID** and **storage_image_ID** parameters.

Invoking the lssi command

dscli> lssi -l

The resulting output

Name	ID	Storage Unit	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 unit.

 \mathbf{ID}^{*}

Indicates the storage image ID that consists of the manufacture, machine type, and serial number.

Storage Unit

Indicates the storage unit ID that consists of the manufacture, machine type, and serial number.

Mode1

Indicates the model number that is associated with the storage unit.

WWNN

Indicates the worldwide node name that is assigned to the storage unit.

State

Indicates the status of the storage unit. One of the following values are displayed:

Online

Indicates that the storage unit is available to process all functions.

Offline

Indicates that the storage unit is offline and not capable of processing any functions.

Resuming

Indicates that the storage unit is preparing to come online.

Quiescing

Indicates that the storage unit is preparing to go offline.

Quiesce Exception

Indicates that the storage unit is in the quiesce exception state.

Forced Quiescing

Indicates that the storage unit is preparing for a force offline operation.

Fenced

Indicates that the storage unit failed and is offline.

Discovery

(DS6000 only) Indicates that the storage unit is determining which physical configurations are available and updates itself when it discovers new hardware.

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 unit.

Key:

- * Displayed when the **-s** parameter is specified.
- + Displayed only when the **-1** 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.

►►_showsi____storage_image_ID____

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	Desc	ID	Storage Unit	Model	WWNN	Signature
My Storage Image	This is my DS storage Image	IBM.2107-75FA120	IBM.2107-75FA120	921	3007ACF3012399E0	0123-4500-0000

State	ESSNet	Volume Group	Os400 Serial	NVS Memory	Cache Memory	Processor Memory
Online	Enabled	-	-	8 GB	128 GB	1 GB

MTS	numeg supported	ETAutoMode	ETMonitor	IOPMmode	ETCCMode	ETHTMMode
IBM.2421 -75FA120	1	tiered	all	Monitored	Enabled	Enabled

Report field definitions

Name

Specifies the name that you assigned to the storage unit.

Desc

Specifies the value that you assigned as a description for the storage unit.

ID Specifies the storage image ID that consists of the manufacture, machine type, and serial number.

Storage Unit

Specifies the storage unit ID that consists of the manufacture, machine type, and serial number.

Mode1

Specifies the model number that is associated with the storage unit.

WWNN

Specifies the worldwide node name that is assigned to the storage unit.

Signature

Specifies the machine signature that is represented by 12 hexadecimal digits in the format *xxxx-xxxx*.

State

Specifies the status of the storage unit. One of the following values are displayed:

Online

Indicates that the storage unit is available to process all functions.

Offline

Indicates that the storage unit is not capable of processing any functions.

Resuming

Indicates that the storage unit is starting to come online.

Quiescing

Indicates that the storage unit is starting to go offline.

Quiesce Exception

Indicates that the storage unit is in the quiesce exception state.

Forced Quiescing

Indicates that the storage unit is starting a force offline operation.

Fenced

Indicates that the storage unit failed and is offline.

Discovery

Indicates that the storage unit is determining which physical configurations are available and updates itself when it discovers new hardware (DS6000 only).

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.

OS400Serial (DS6000 only)

Specifies " - " for a DS8000 model and the iSeries serial number for a DS6000 model.

The serial number consists of three hexadecimal characters. It is used to uniquely identify LUNs within a customer's storage complex. It is appended to the unitSerialNumber that is returned by a SCSI inquiry command that is directed to each LUN.

NVS Memory

Specifies the amount in decimal gigabytes (GB) of nonvolatile storage (NVS) memory that is configured on the storage unit. Example: 4.0 GB

Cache Memory

Specifies the amount in decimal gigabytes (GB) of cache memory that is configured on the storage unit. This memory equals the processor memory less the memory that is required for the operating system (varies with microcode 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 unit. This memory equals the installed memory less the reserved memory requirements, including firmware memory, and varies with the DS8000 model, the DS8000 configuration, and the microcode level.

MTS

Specifies the order type of the storage unit. The order type and the machine type of the storage unit is the same on all storage units that are ordered before release 2.4. After release 2.4, the order type varies according to the warranty periods that are associated with the storage unit.

Note: This value is not reported for a DS6000 model. A " - " value is displayed.

numegsupported

Specifies the number of encryption 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

Indicates whether volumes are monitored by the IBM Easy Tier LIC feature.

all Easy Tier monitors all of the volumes on the DS8000 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 DS8000 system.

"-" The Easy Tier feature is not supported.

IOPMmode

Specifies the I/O priority management mode, which can include one of the following values:

Disabled

Specifies that the I/O priority manager function is disabled.

Monitored

Specifies that resources associated with performance groups that specify I/O management are monitored, but not managed.

MonitoredSNMP

Specifies the same condition as Monitored, but with added SNMP trap support.

Managed

Specifies that resources associated with performance groups that specify I/O management are managed.

ManagedSNMP

Specifies the same condition as Managed, but with added SNMP trap support.

Note: The default I/O priority management mode is managed mode.

ETCCMode

Easy Tier cooperative caching. The Easy Tier cooperative caching function is also known as Easy Tier Server.

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 System Storage 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 DS8000 system does not support the Easy Tier heat map transfer feature.

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.

setioport


Configures one or more I/O ports for open systems or z Systems host system connections.

showioport

Displays the properties of a specified I/O port. It optionally displays the performance metrics for the I/O port.

Isioport

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, 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 **-1** and **-s** parameters together.

-1

(Optional) Displays default output plus the I/O port interface speed.

-type fc | escon

(Optional) Displays I/O ports of the specified port type. *Escon* can be specified as the port type for DS8000 models only.

-topology fc-al | scsi-fcp | ficon

(Optional) Displays Fibre Channel I/O ports with the specified topology.

-state online | offline | fenced | deconfigured

(Optional) Displays I/O ports of the specified state.

-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, 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:

For **DS8000**:

- *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).

For **DS6000**:

- *EE* is an I/O port enclosure number in the range of 00 01.
- *A* is the adapter number and is specified as 0, 1, 2, or 3.
- *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.

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 reports that are associated with the **lsioport** command.

Invoking the lsioport command

dscli> lsioport -dev IBM.2107-75FA120 -1

The resulting output

ID	WWPN	State	Туре	Торо	Portgrp	Speed
IBM.2107 -75FA120 /I0111	307BCF30 A3299E0	Online	Fibre Chan- nel-LW	SCSI-FCP	0	1Gb/s

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:

Online

Indicates that the storage system can process all functions (default).

Offline

Indicates that the storage system is not capable of processing any functions.

Resuming

Indicates that the storage system is in the process of coming online.

Quiescing

Indicates that the storage system is in the process of going offline.

Fenced

Indicates that the storage system failed and is offline.

Deconfigured

Indicates that the I/O port is being deleted.

Туре

Indicates 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

Торо

Indicates the I/O port topology. The following values can be displayed:

- FC-AL
- 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
- FICON ports = 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s, 16 Gb/s
- Unknown

Report field definitions (with the -metrics parameter)

When you use the **-metrics** parameter and do not specify a port type, two reports are displayed.

- For DS8000 models, one report is for the FICON/ESCON I/O port type and the other report is for the SCSI-FCP I/O port type. A banner is displayed (for example: ===FICON/ESCON-Enabled I/O Ports===) before each report.
- For DS6000 models, one report is for the FICON I/O port type and the other report is for the SCSI-FCP I/O port type. A banner is displayed (for example: ===FICON-Enabled I/O Ports===) before each report.

Note: For DS8000 and DS6000 models, a report is not displayed for a port type that has no enabled ports.

FICON ports (for DS6000 models only) and **FICON/ESCON ports** (for DS8000 models only): Each of the following headers and value types are 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 (ECKDTM) 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 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 are 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 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.

Report field definitions (with -metrics and -I parameters)

Fibre Channel Link Errors: Each of the following headers and value types are 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 has changed from an active (AC) state to a Link Recovery (LR1) state.

LRRec

This count is the number of times the port has 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 has been detected as previously processed.

InvRel0ffset

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 has 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 are 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 one of three 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 1310nm port.

Laser LL 1550-LW

Specifies a long-wave laser LL 1550nm port.

Non-Optical

Specifies that the port is not optical or is of another type.

Unknown

Specifies that the port type could not be 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

Note: A dash (-) specifies that data was not available.

FECStatus

Indicates the forward error correction (FEC) status on the link:

Active

Specifies that the forward error correction is active on the link.

Inactive

Specifies that forward error correction is not active on the link.

UncorrectedBitErr

Indicates the number of bad data blocks that were not corrected by forward error correction.

Note: A dash (-) means that the data was not available.


CorrectedBitErr

Indicates the number of bad data blocks that were corrected by forward error correction.

Note: A dash (-) means that the data was not available.

setioport

The **setioport** command configures one or more I/O ports for open systems or System 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 will temporarily override any defined value for *devid* for the current command.

-topology fc-al | scsi-fcp | ficon

(Optional) Sets the topology for an I/O port, either Fibre Channel Arbitrated Loop, SCSI-FCP, or FICON.

fibre channel arbitrated loop (code *fc-al*)

The fc-al topology setting enables the SCSI ULP with a FC-AL topology. The FC-AL topology does not support PPRC path I/O operations nor 16 Gb/s host adapters.

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.

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:

For **DS8000**:

- *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).

For **DS6000**:

- *EE* is an I/O port enclosure number in the range of 00 01.
- *A* is the adapter number and is specified as 0, 1, 2, or 3.
- *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.

Example

Invoking the setioport command

This example configures four I/O ports for FICON topology. dscli> setioport -dev IBM.2107-75FA120 -topology ficon I0111 I0121 I0211 I0221

The resulting output

I/O Port I0111 successfully configured.I/O Port I0121 successfully configured.I/O Port I0211 successfully configured.I/O Port I0221 successfully configured.

showioport

The **showioport** command displays properties of an I/O port. It optionally displays the performance metrics for a specific I/O port.

```
_____showioport_______port_ID______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 temporarily overrides any defined value for *devid* for the current command.

-metrics

(Optional) Specifies that the port ID and the performance metrics for the specified I/O port be displayed.

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 | -

(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:

DS8000

- *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).

DS6000:

- *EE* is an I/O port enclosure number in the range of 00 01.
- *A* is the adapter number and is specified as 0, 1, 2, or 3.
- *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

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 **showioport** command.

Invoking the showioport command to show port properties

dscli> showioport -dev IBM.1400-1B1 I0000

The resulting output

ID	WWPN	State	Loc	Туре
IBM.1400 -1B1/ I0000	50050076 3080005AE	Online	U1400. 1B1. 1300025-P1- C1-T0	Fibre Channel-SW

Speed	Торо	Portgrp	unkSCSIlog	physloc
8 Gb/s	SCSI-FCP	0	-	R1-I1-C1-T0

Report field definitions (with the -dev 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:

Online

Indicates that the storage system can process all functions (default).

Offline

Indicates that the storage system is not capable of processing any functions.

Resuming

Indicates that the storage system is in the process of coming online.

Quiescing

Indicates that the storage system is in the process of going offline.

Fenced

Indicates that the storage system failed and is offline.

Deconfigured

Indicates that the I/O port is being deleted.

Loc

Specifies the storage enclosure location by identifying the storage unit frame that contains the storage enclosure. The location format is *Utttt.mmm.ppsssss*.

Туре

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
- FICON ports = 1 Gb/s, 2 Gb/s, 4 Gb/s, 8 Gb/s , 16 Gb/s
- Unknown

Торе

Specifies the port topology. If the port type is not Fibre Channel, then the displayed value is " - ". One of the following values is displayed:

- FC-AL
- 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.

unkSCSI1og

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

This information is not supported for the DS6000.

Note: R1-I3-C2-P1 is an example of this format.

Example 2

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 **showioport** command, with the -metrics parameter.

Invoking the showioport command to show port performance metrics

dscli> showioport -dev IBM.1400-1B1 -metrics I0000

The resulting output

ID	Date	byteread (FICON/ESCON)	bytewrit (FICON/ESCON)	Reads (FICON/ESCON)
IBM.1400 -1B1/ I0000	04/17/2012 11:02:45 MST	0	0	0

Writes (FICON/ESCON)	timeread (FICON/ESCON)	timewrite (FICON/ESCON)	bytewrit (PPRC)	byteread (PPRC)
0	0	0	0	0

Writes	timeread	timewrite	CmdRetries (FICON)	TransferReady
(FICON/ESCON)	(FICON/ESCON)	(FICON/ESCON)		(FICON)
0	0	0	0	0

bytewrit (PPRC)	byteread (PPRC)
0	0

Writes (PPRC)	Reads (PPRC)	timewrite (PPRC)	timeread (PPRC)	byteread (SCSI)
0	0	0	0	1271326535

bytewrit (SCSI)	Reads (SCSI)	Writes (SCSI)	timeread (SCSI)	timewrite (SCSI)
65721755	2479310687	1170136080	81426674	407813158

LinkFailErr (FC)	LossSyncErr (FC)	LossSigErr (FC)	PrimSeqErr (FC)	InvTxWordErr (FC)
8	45	0	0	125

CRCErr (FC)	LRSent (FC)	LRRec (FC)	IllegalFrame (FC)	OutOrdData (FC)
0	0	0	0	0

OutOrdACK (FC)	DupFrame (FC)	InvRelOffset (FC)	SeqTimeout (FC)	BitErrRate (FC)
0	0	0	0	0

RcvBufZero (FC)	SndBufZero (FC)	RetQFullBusy (FC)	ExchOverrun (FC)	ExchCntHigh (FC)
0	0	0	0	0

ExchRemAbort (FC)	CurrentSpeed (FC)	%UtilizeCPU (FC)
0	0	0

TxPower (RDP)	RxPower (RDP)	TransceiverTemp (RDP)	SupplyVolt (RDP)	TxBias (RDP)
-2.8 dBm (525.2 uW)	-3.6 dBm (440.5 uW)	65 C	3304.7 mV	8.484 mA

ConnectorType (RDP)	TxType (RDP)	FECStatus(RDP)	UncorrectedBitErr (RDP)	CorrectedBitErr (RDP)
SFP+	Laser-SW	Active	315	250

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 requested for transport mode write operations because not enough buffers were available to receive unsolicited data.

Note: A dash (-) means that the data was 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 was 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 receiving 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 was 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 was 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 was 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 was 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 was 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 could not be 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.

UncorrectedBitErr (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 was not available.

CorrectedBitErr (RDP)

Specifies the number of bad data blocks that were corrected by forward error correction (FEC).

Note: A dash (-) means that the data was not available.

Host connection commands

Various commands are available to configure host connections and to display host connection information.

The following host connection commands are available:

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.

1shostconnect

Displays a list of host connections for a storage image and the status information for each host connection in the list.

lshostvol

Displays the mapping of host device names or volume names to machine type 2105, 2107/242x, and 1750 volume IDs.

1sportprof

Displays a list of port profiles that are supported on a storage unit 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.

mkhostconnect

Configures the open systems hosts port attachments to Fibre Channel ports that are configured for FC-AL or SCSI-FCP topology.

rmhostconnect

Removes a SCSI host port connection from a storage image.

showhostconnect

Displays the detailed properties of a specified storage image host connection.


lshosttype

Displays a list of known hosts, their associated port profiles, address discovery, and logical block size values.

chhostconnect

The chhostconnect command modifies a SCSI host port configuration.

```
►→—chhostconnect
```



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 unit. 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 issue this command allows you to save input and processing time. The **-hosttype** parameter supplies the same information as if you had used the following three parameters:
 - -profile
 - -addrdiscovery
 - -1bs
- **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, 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.

For DS8000, example of a fully qualified storage image ID: IBM.2107-75FA120

-1bs 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 only applies 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.

-hosttype host_type

(Optional) Specifies information about the following three parameters:

- -profile
- -addrdiscovery
- -1bs

Notes:

- 1. You cannot use the **-hosttype** parameter with 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 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 have not been 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. You can list one or more port IDs separated by commas with no spaces. If you enter a list of I/O port IDs, access from the specified host connection to the specified volume group is allowed using only the designated list of port IDs.

- **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 unit is configured to allow host connection access to the specified volume group using any I/O port that is configured for FC-AL or 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:

• For DS8000,

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:

• For DS8000:

dscli> chhostconnect -ioport IBM.2107-75FA120/I0222 IBM.2107-75FA120/1

where *IBM.1750-68FA120/1* or *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:

For DS8000:

- 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).

For DS6000:

- *EE* is an I/O port enclosure number in the range of 00 01.
- *A* is the adapter number and is specified as 0, 1, 2, or 3.
- *P* is the port number (0 3).

To specify a range of port IDs, separate the port IDs with a hyphen.

Separate multiple port IDs or ranges of port IDs with a comma between each ID or range of IDs.

Note: Changing the I/O port values can result in a disruption of current logins by the host systems.

-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.

-name new_host_name

(Optional) Specifies the user-assigned host system or port name. The name is limited to 32-byte or 16 double-byte characters.

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:

• For DS8000:

dscli> chhostconnect -desc newdescription IBM.2107-75FA120/1

• For DS6000:

dscli> chhostconnect -desc newdescription IBM.1750-68FA120/1

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 have 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.

Ishostconnect

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 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.

For DS8000, example of a fully qualified storage image ID: IBM.2107-75FA120

- S

(Optional) Specifies the host connection IDs. You cannot use the -1 and the -s parameters together.

-1

(Optional) Specifies the default output and your description for each host connection in the list. You cannot use the **-1** 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 unit, be displayed.

Note: The list of logged in host posts includes all of the host ports that the storage unit detects, and it does not take into account changes that the storage unit could not detect. For example, the storage unit cannot detect that a cable has been disconnected from the port of the host device or that a fabric zoning change has 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 for which you have created for host connection, as well as 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 is 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.

For DS8000, 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 -1

The resulting output The following table displays four ports, which is possible when you query a DS8000 model. If you query a DS6000 model, only two ports are displayed.

Name	ID	WWPN	HostType	LBS	Addrdis covery
My host port 1	IBM.2107- 75FA120/1	3007ACF3 0A2399E0	Unknown	512	reportLUN
My host port 2	IBM.2107- 75FA120/2	3007ACF3 0A2399E1	Unknown	512	reportLUN
My host port 3	IBM.2107- 75FA120/3	3007ACF3 0A2399E2	Unknown	512	reportLUN
My host port 4	IBM.2107- 75FA120/4	3007ACF3 0A2399E3	Unknown	512	reportLUN

Profile	portgrp	volgrpID	atchtopo	ESSIOport	Speed	Desc	host
IBM pSeries - AIX	0	V100	-	I0111,I0121 I0211,I0221	Unknown	SCSI1	Production01
IBM pSeries - AIX	0	V100	-	All	Unknown	SCSI2	Production01
IBM pSeries – pLinux	0	V100	-	I0111,I0121 I0211,I0221	Unknown	SCSI3	
IBM pSeries - pLinux	0	V100	-	-	Unknown	SCSI4	

Example of Ishostconnect using the -unknown parameter

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

WWNN	WWPN	ESSIOport
3007ACF30A239900	3007ACF30A2399E1	I0121
3007ACF30A239900	3007ACF30A2399E2	I0111, I0121, I0211, I0221
3007ACF30A239900	3007ACF30A2399E3	I0111

Example of Ishostconnect using the -login parameter

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 unit.

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 System i host system attachments.

Addrdiscovery

Indicates the LUN address discovery method 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.

Indicates the volume group ID. This ID is a unique identifier within the DS6000 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.

ESSI0port

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:

For **DS8000**:

- *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).

For **DS6000**:

- *EE* is an I/O port enclosure number in the range of 00 01.
- *A* is the adapter number and is specified as 0, 1, 2, or 3.
- *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 displays 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.

ESSI0port

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:

For **DS8000**:

- *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).

For **DS6000**:

- *EE* is an I/O port enclosure number in the range of 00 01.
- *A* is the adapter number and is specified as 0, 1, 2, or 3.
- *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.

ESSI0port

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:

For **DS8000**:

- *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).

For **DS6000**:

- *EE* is an I/O port enclosure number in the range of 00 01.
- *A* is the adapter number and is specified as 0, 1, 2, or 3.
- *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 unit.

Ishostvol

The **lshostvol** command displays the mapping of host device names or volume names to machine type 2105, 2107, and 1750 volume IDs. (This command is not supported on the i5/OS.)

► lshostvol ----

Parameters

There are no parameters for this command.

Notes:

- 1. The **lshostvol** command displays only volumes that are accessed 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 only has 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 executed by the mass storage subsystem. To initiate an I/O request, the creates a small data structure in memory, appends it to a "send" queue, and if this 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 at this time even though it has both MSCP and direct paths.
- 2. If you do not have installed 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 does not occur automatically, you must issue the following 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 will return 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 have SDD installed. 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.2107-75FA130/175F	vpath03
my_vol_07,my_vol_09	IBM.2105-29239/119E	Vpath04
my_vol_08,my_vol_10	IBM.2105-29239/119F	Vpath05

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	-
Hdisk07	IBM.2105-29239/119E	-

Device/Volume Name	Volume ID	Vpath Name
Hdisk08	IBM.2105-29239/119F	-

Report column definitions

Device/Volume name

Specifies the nickname 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 unit.

Vpath name

Specifies the virtual path name. When SDD is not installed, this value is reported as " - ".

Isportprof

The **lsportprof** command displays a list of port profiles that are supported on a storage unit 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.

```
► lsportprof storage_image_ID
```

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 given 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.

managehostconnect

The managehostconnect command modifies the volume group assignment for a SCSI host port.

▶ → — managehostconnect-	 -volgrp	>
▶ port_grp_number		

Parameters

Notes:

- 1. The **managehostconnect** command can be disruptive to host system I/O operations if the affected host port is logged onto 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 have issued the **lshostconnect** or **showhostconnect** commands and have analyzed 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 will temporarily override 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 will be modified to access this volume group ID. A volume group ID is a four-digit decimal number with no leading zeroes, 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.

mkhostconnect

The **mkhostconnect** command configures open systems hosts port attachments to Fibre Channel ports that are configured for FC-AL or SCSI-FCP topology.



Parameters

Open systems hosts port attachments to Fibre Channel ports are configured for identified access mode and SCSI protocol.

Notes:

- Ensure that you use the -hosttype parameter when you issue this command, because doing so saves input and processing time. The -hosttype parameter supplies the same information as if you had used the following three parameters:
 - -profile
 - -addrdiscovery

```
• -1bs
```

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 will temporarily override 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: You should not create more than one hostconnect per WWPN, except for SUN hosts. Creating more than one hostconnect per WWPN (each with a different volume group) is only supported for SUN.

-1bs 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 System i models (OS/400).

Notes:

- 1. Do not use the **-1bs** 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 there are additional considerations 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 only applies 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
- -1bs

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 have not been 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 zeroes, 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 issue 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 the same WWPN. Ensure that there are no conflicts 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. You can list one or more port IDs separated by commas with no spaces. If you enter a list of I/O port IDs, access from the specified host connection to the specified volume group is allowed using only the designated list of port IDs.

- **all** Specifies that you want to add all I/O port IDs. This 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 unit is configured to allow host connection access to the specified volume group using any I/O port that is configured for FC-AL or SCSI-FCP topology.

A port ID is four hexadecimal characters in the format *EEAP*, where: For DS8000:

- *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).

For DS6000:

- *EE* is an I/O port enclosure number in the range of 00 01.
- *A* is the adapter number and is specified as 0, 1, 2, or 3.
- "*P*" is the port number (0 3).

This number is prefixed with the letter *I*.

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.

-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_name | -

(Required) Specifies your host system or port name, limited to 16 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 -dev IBM.2107-75FA120 -wwname 12341234000A000F
-profile "IBM pSeries - AIX" host 1 port 1
```

The resulting output

Host connection 0 successfully created.

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 will temporarily override 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

For DS6000, example of a fully qualified host connection ID: IBM.1750-68FA120/2

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 have 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.

showhostconnect

The showhostconnect command displays detailed properties of a storage image host connection.

► showhostconnect _____ host_connection_ID ____ host_connection_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 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	Production01

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 unit.

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. This 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.

Indicates the volume group ID. This ID is a unique identifier within the DS6000 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.

ESSI0port

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:

For **DS8000**:

- *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).

For **DS6000**:

• *EE* is an I/O port enclosure number in the range of 00 - 01.

- *A* is the adapter number and is specified as 0, 1, 2, or 3.
- *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 displays in the **Name** column. If a name is not available, a value of "-" displays.

Ishosttype

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.

►►—lshosttype—	-type— volumeGroup_type—	>
└── -S── └── -1──		

Parameters

- S

(Optional) Displays the host types only. You cannot use the -1 and -s parameters together.

-1

(Optional) Displays the default output for the specified host type. You cannot use the **-1** 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 given 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 64K 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 additional host type details.

Key:

- * Displayed when the **-s** parameter is specified.
- + Displayed only when the **-1** parameter is specified.

Storage configuration commands

There are some specific DS CLI commands that are associated with configuring storage for z Systems and open system hosts.

The commands below allow you to configure storage for z Systems 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.

Isarraysite

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 -1 and the -s parameters together.

-1

(Optional) Displays the default output plus the disk drive module rpm (revolutions per minute), disk class, and encryption support capability. You cannot use the **-1** 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 | n1 | ent | sata | flash

(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.
- flash Specifies high-performance flash devices.

-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

Arsite	DA pair	Dkcap (10^9B)	Diskrpm	State	Array	Diskclass	Encrypt
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.

Flash

Indicates high-performance flash devices.

NL Indicates near-line and represents ATA (FATA) disk drives

SATA

Indicates high capacity SATA disk drives.

SSD

Indicates solid-state devices.

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 -1 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. IBM DS6000 array sites consist of four DDMs. IBM DS8000 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. One of the following values are 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.
- Flash Indicates high-performance flash devices.
- NL Indicates near-line and represents ATA (FATA) disk drives.
- SATA Indicates high capacity SATA disk drives.
- SSD Indicates solid-state devices.

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 unit.


showarray

Generates a report that displays the detailed properties of a specific array.

Isarray

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 -1 and the -s parameters together.

-1

(Optional) Displays default output plus the disk class and encryption capability of the arrays. You cannot use the **-1** 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 like: For DS8000,

dscli> lsarray IBM.2107-75FA120 -1 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 -1

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)

Arsite	Rank	DA Pair	DDMcap (10^9B)	Diskclass	Encrypt
S20	R11	10	145	ENT	supported
S21	R12	11	145	ENT	supported
S30	R13	20	300	NL	unsupported
S31	-	21	300	SATA	unsupported

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 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.

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
- Flash Indicates high-performance flash devices.
- NL Indicates near-line and represents ATA (FATA) disk drives

SATA Indicates high capacity SATA disk drives.

SSD Indicates solid-state devices.

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 **-1** parameter is specified.

mkarray

The **mkarray** command creates one array per command. You can specify two array sites if you are working with a DS6000 machine type, but 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.

-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

(Required for a DS6000 machine type) Specify one or two array sites for IBM 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.

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 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 (10⁹B) 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.
- Flash Indicates high-performance flash devices.
- **NL** Indicates near-line and represents ATA (FATA) disk drives.
- **SATA** Indicates high capacity SATA disk drives.
- SSD Indicates solid-state devices.

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 unit 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 unit.

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 a 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 will also stop the rank depopulation.

Note: Versions of the DS CLI prior to Release 6.1 will fail if they use the *-reserve* parameter on Release 6.1 or later DS8000 machines 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 there are any extents on this rank that are used to provision a volume, the unassign will normally fail. 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 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 will be given, but may be suppressed with the **-quiet** parameter.
- Versions of the DS CLI prior to Release 6.1 will fail if they use the -unassign parameter on Release 6.1 or later DS8000 machines with allocated extents. These previous versions of the DS CLI must manually delete or migrate any volumes using these ranks before using 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.

Isrank

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.

For DS8000, example: IBM.2107-75FA120

- S

(Optional) Displays only the rank ID. You cannot use the -1 and the -s parameters together.

-1

(Optional) Displays the default output plus the extent pool name, number of extents, extents that are used, and encryption group for each rank. You cannot use the **-1** 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 | configpending | configpendingerr | configoutofsyncerr
```

(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.

-encryptgrp encryption_group_ID

(Optional) Displays the encryption group that this rank uses.

-marray managed_array_ID

(Optional) Displays only those ranks that are assigned to the specified managed array.

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

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 **lsrank** command by using the **-l** parameter.

Invoking the Isrank command

dscli>lsrank -dev IBM.2107-75FA120 -1

The resulting output

ID	Group	State	Datastate	Array	RAIDtype
IBM.2107- 75FA120/R1	0	Normal	Normal	A1	5
IBM.2107- 75FA120/R2	0	Normal	Normal	A2	5
IBM.2107- 75FA120/R3	0	Normal	Normal	A3	5
IBM.2107- 75FA120/R4	0	Normal	Normal	A4	5

ExtpoolID	Extpoolnam	Stgtype	Exts	Usedexts	Encryptgrp	marray
P1	poolA	fb	1,000	500	1	MA1
P1	poolA	fb	1,000	500	1	MA2
P1	poolA	fb	1,000	500	1	MA3
P1	poolA	fb	1,000	500	1	MA4

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 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 the depopulation of a rank failed and efforts to depopulate the rank stopped. 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 failed to complete successfully. This state reflects an internal error condition and not an error in the request to remove the rank. To correct this state, you must reissue the **rmrank** command for the designated rank configuration.

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. The value that is displayed is a number in the range of 1 - 4000.

Usedexts

Indicates the number of extents that are allocated to volumes from the designated rank. The value that is displayed is a number in the range of 1 - 4000.

encryptgrp

Indicates the encryption group number. 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. 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 will be 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*.

mkrank

The **mkrank** command creates one fixed block or count key data (CKD) rank from one array.



Parameters

Notes:

- 1. The DS8000 system assigns the ID 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 when using the **mkrank** command. Using either of these parameters allows you to be notified if the rank configuration fails for any reason.

-dev storage_image_ID

(Optional) Specifies the storage image ID, which includes manufacturer, type, and serial number. The storage image ID is required if you do not specify fully qualified IDs 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 will temporarily override 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 zeroes, prefixed with the letter A.

-stgtype fb | ckd

(Required) Specifies the type of extent for which the rank will be configured, either fixed block or count key data.

-encryptgrp encryption_group_ID

(Optional) Specifies the encryption group that this rank should use. The default is zero, which means that no encryption 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, then the rank will be assigned to the extent pool. Otherwise, the rank state is unassigned. If specified, the extent pool 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 zeroes, prefixed with the letter *P*.

Note: You must use the chrank command if you choose to specify the extent pool ID at a later time.

Example

Invoking the mkrank command

```
dscli> mkrank -dev IBM.2107-75FA120
-array A44 -stgtype fb -wait -encryptgrp 1
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Device: IBM.2107-75FA120

Rank IBM.2107-75FA120/R44 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 (for DS8000) or storage unit (for DS6000) 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.

►→—showrank______ rank_ID____ rank_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 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

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 **showrank** command.

Invoking the showrank command to show rank properties

dscli> showrank -dev IBM.2107-75FA120 R34

The resulting output

ID	SN	Group	State	Datastate	Array	RAIDtype
2107 -75FA120 /R34	A23567	0	Normal	Normal	IBM.2107 -75FA120 /A44	10

ExtpoolID	Extpoolnam	Volumes	Stgtype	Exts	Usedexts
IBM.2107 -75FA120 /P48	host_4 _extpool	IBM.2107 -75FA120 /R7	FB	1,000	500

Widearrays	Nararrays	Trksize	Strpsize	Strpesize	Extsize	Encryptgrp
1	0	128	4	4	16,384	-

migrating(in)	migrating(out)
0	10

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 the depopulation of a rank has failed and efforts to depopulate the rank have stopped. 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. 1 - 4000 are valid values.

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 a 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.

encryptgrp

Indicates the encryption 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 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 migrating out of the rank. In other words, this value is the number of migrating extents in the rank that have specified this rank as the source of the migration.

Example 2

Invoking the showrank command to show performance metrics

dscli> showrank -dev IBM.2107-75FA120 -metrics R34

The resulting output

ID	Date	Byteread	Bytewrit	Reads
2107 -75FA120 /R34	10/11/04 02:23:47	10000	10000	10000

Writes	Timeread	Timewrite	dataencrypted
10000	10000	10000	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 unit 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 unit.

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

Note: The -virextentlimit, -virlimit, and -virthreshold parameters are used only on DS8000 models.

-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.

Note: A warning is generated when this percentage is exceeded by capacity allocation to Extent Space Efficient (ESE) volumes.

-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.

Isextpool

The **lsextpool** command displays a list of extent pools in a storage unit and status information on each extent pool 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 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 extent pool IDs. You cannot use the -1 and the -s parameters together.

-1

(Optional) Displays default output plus the number of ranks and tiers, encryption group ID, and management status of the extent pools. You cannot use the **-1** 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.

-encryptgrp encryption_group_ID

(Optional) Displays only extent pool of the specified encryption 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

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 **lsextpool** command using the **-l** parameter.

Invoking the lsextpool command

dscli> lsextpool -dev IBM.2107-75FA120 -1

The resulting output

Name	ID	Stgtype	Rankgrp	Status	Avail- stor (2^30B)
ckd_c0 _ext _pool00	IBM.2107- 1300321 /P0	ckd	0	below	600
ckd_c1 _ext _pool01	IBM.2107- 1300321 /P1	ckd	1	below	600
fb_c0 _ext _pool02	IBM.2107- 1300321 /P2	fb	0	below	715
fb_c1 _ext _pool03	IBM.2107- 1300321 /P3	fb	1	below	715

%allo- cated	Avail- able	Reser- ved	Num- vols	Num- ranks	encrypt grp
21	681	0	64	1	-
21	681	0	64	1	-
8	715	0	64	1	-
8	715	0	64	1	-

numtiers	etmanaged
1	no

Report field definitions

Name

Indicates the name you assigned to the extent pool.

\mathbf{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 %Extents Available is 0.

Availstor (2³⁰B)

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.

Encryptgrp⁺

The encryption 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.

Numtiers⁺

Indicates the number of tiers in this extent pool.

$\mathsf{Etmanaged}^{+}$


Indicates whether or not the extent pool is managed.

Key:

- * Displayed when the **-s** parameter is specified.
- + Displayed only when the **-1** 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.

-extpool 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 DS8000 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. The -virextentlimit, -virlimit, and -virthreshold parameters are used only on DS8000 models.
- **3**. An extent pool object is assigned to either rank group 0 or 1, which allows the extent pool to be managed by storage unit server 0 or 1 respectively.
- 4. Create extent pool objects before creating array and rank objects.
- 5. Create extent pools of a given type for both rank groups 0 and 1 so 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 encryption 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.

-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 so 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 sent and the real extent status is reported as exceeded.

-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.

-encryptgrp encryption_group_ID

(Optional) Specifies the encryption group that this extent pool should use. The default is zero, which means that no encryption 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 -encryptgrp 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

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 **showextpool** command.

Invoking the showextpool command to show extent pool properties

dscli> showextpool -dev IBM.2107-75FA120 P21

The resulting output

Name	ID	stgtype	totlstor (2^30B)	availstor (2^30B)	resvdstor (2^30B)	rankgrp
host_4 _extpool	IBM.2107- 75FA120 /P21	fb	1000	800	0	1

num ranks	numvols	status	%allo- cated	%avail- able	config- ured	allowed	avail- able
4	3	exceeded	20	80	1000	600	800

allocated	reserved	%limit	%thres- hold
200	0	80	30

virextstatus	%virallocated	%viravailable	virconfigured	virallowed
below	10	90	1000	800

viravailable	virallocated	virreserved	%virextlimit	%virext- threshold	encrypt grp
720	80	0	70	35	1

%allocated	%allocated	%allocated	%allocated
(ese)	(rep)	(std)	(over)
0	0	0	0

%virallocated	%virallocated	%virallocated
(ese)	(tse)	(init)
3	3	0

%migrating(in)	%migrating(out)
0	10

numtiers	etmanaged
1	no

etmigpauseremain	etmonpauseremain	etmonitorreset
-	-	2013-07-26T14:00:00+07

Report field definitions

Name

Specifies the name 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 (2³⁰ 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 (2³⁰ 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.

%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.

Virextstatus

(DS8000) 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.

bel ow

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.

(DS6000) No value displayed.

%virallocated

(DS8000) Specifies the percentage of virtual extents that are allocated compared to the total virtual extents that are allowed. Valid values are 0 - 100.

(DS6000) No value displayed.

%viravailable

(DS8000) Specifies the percentage of virtual extents that are available compared to the total virtual extents that are allowed. Valid values are 0 - 100.

(DS6000) No value displayed.

Virconfigured

(DS8000) Specifies the number of virtual extents that are configured in the extent pool.

Virallowed

(DS8000) Specifies the number of virtual extents that are below the applicable virtual extent limit.

(DS6000) No value displayed.

Viravailable

(DS8000) Specifies the number of virtual extents that are available for allocation to space-efficient volumes.

(DS6000) No value displayed.

Virallocated

(DS8000) Specifies the number of virtual extents in the extent pool that are allocated to space-efficient volumes.

(DS6000) No value displayed.

Virreserved

(DS8000) 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.

(DS6000) No value displayed.

%virextlimit

(DS8000) 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.

(DS6000) No value displayed.

%virextthreshold

(DS8000) 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.

(DS6000) No value displayed.

encryptgrp

The encryption 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.

%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:

уууу	the year			
MM	the month (01-12)			
dd	the day (01-31)			
Т	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.

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

(DS8000) Specifies the virtual extent pool capacity in gibibytes (GiB).

(DS6000) No value displayed.

virext

(DS8000) Specifies the number of virtual extents in the extent pool.

(DS6000) No value displayed.

dataencrypted

Specifies whether the data stored on the physical media is encrypted.

DSCLI

Example 3

Invoking the showextpool command with the -tier option

0	I
dscli> showextpool ·	-tier p2
	, 2013 10:44:07 AM MST IBM
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
%limit	100
%threshold	15
virextstatus	full
%virallocated	100
%viravailable	0
virconfigured	550
virallowed	550
viravailable	0
virallocated	550
virreserved	0
%virextlimit	-
%virextthreshold	-
encryptgrp	-
%allocated(ese)	15
%allocated(rep)	0
%allocated(std)	0
%allocated(over)	0
%virallocated(ese)	100
%virallocated(tse)	0
<pre>%virallocated(init)</pre>	0
<pre>%migrating(in)</pre>	0

%migrating(out)	0			
numtiers	2			
etmanaged	yes			
etmigpauseremain	-			
etmonpauseremain	-			
etmonitorreset 2013-07-26T14:00:00+07				
========Tier Distribution=========				
Tier Cap (GiB/mod1)	%allocated %	kassigned		
SSD 2121	84	Θ		
ENT 436981	100	14		
dscli>				

Report field definitions

Tier ID

SSD Solid state device tier.

- **ENT** Enterprise tier; consists of drives with speeds of 10K RPM, 15K RPM, or a combination of drives of both speeds.
- NL Near Line tier; consists of high volume SATA or SAS Near Line disk drives.

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.

Address group specific commands

A specific DS8 CLI command is used to display address group information.


The following address group command is available:

1saddressgrp

This command generates a report that displays a list of address groups for a storage unit. It also provides the status information for each address group in the list.

Isaddressgrp

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.

For DS8000, example: IBM.2107-75FA120

For DS6000, example: IBM.1750-68FA120

- -s (Optional). Displays the address group IDs only. You cannot use the -1 and the -s parameters together.
- -1 (Optional). Displays the default output. You cannot use the -1 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 -1

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

\mathbf{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 -1 parameter is specified.

Logical control unit specific commands

Commands are referenced for tasks that are associated with System z logical control units.

Use the following logical control unit specific commands:

chlcu Modifies a logical control unit.

- **1slcu** Generates a report that displays a list of logical control units for a storage unit and the status information for each logical control unit in the list.
- **mklcu** Creates a logical control unit in a storage unit. A logical control unit is configured to represent a grouping of logical CKD volumes.


rmlcu Deletes one or more specified logical control units.

show1cu

Generates a report that displays the detailed properties for the specified logical control unit.

chicu

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.

-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 for DS8000 and 00 - 1F for DS6000.

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.
```

lsicu

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 unit that is specified. The storage image ID is required if you do not specify a fully qualified LCU 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.

For DS8000, example: IBM.2107-75FA120

-s

(Optional). Displays LCU IDs only. You cannot use the -1 and the -s parameters together.

-1

(Optional). Use this parameter to display the default output plus resource groups. You cannot use the **-1** 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. A LCU ID is two hexadecimal characters 00 - FE for DS8000 and 00 - 1F for DS6000.

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 using the **-l** parameter.

Invoking the lslcu command

dscli> lslcu -dev IBM.2107-75FA120 -1

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.

This is a unique identifier that is assigned to this logical control unit object. The ID includes the storage image ID and the 2-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 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 **-1** 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. The DS6000 has a 16 384 volume address space that is partitioned into 64 logical subsystem (LSS) units, where each LSS contains 256 logical volume numbers. The 64 LSS units are assigned to one of 4 address groups, where each address group contains 16 LSSs, or 4096 volume addresses. All of the LSSs in one address group must be of the same type (CKD or fixed block).
- 5. 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 if you do not specify a fully qualified LCU 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.

For DS8000, example: IBM.2107-75FA120

-qty quantity

(Required). Specifies the number of LCU IDs to be created. You can specify a *quantity* value from 1 - 255 for the DS8000 model. You can specify a *quantity* value from 1 - 64 for the DS6000 model.

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 for DS8000 is 00 - FE and for DS6000 it is 00 - 1F.

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. A LCU ID is two hexadecimal characters 00 - FE for DS8000 and 00 - 1F for DS6000.

Example: 00

-ss ss_ID

(Required). Specifies the subsystem ID that you have assigned. 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 be enabled. 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 additional 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 has been 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 has been received. Otherwise, the volume does not become long busy.

-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 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 suspending a concurrent copy session. 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 suspending the XRC session. 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.

For DS8000, 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 for the DS8000 and 00 - 1F for the DS6000.

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-OF ? 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.

For DS8000, 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 (DS8000 only); 00 - 1F (DS6000 only).

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: May 22, 2015 8:43:04 AM MST IBM DSCLI Version: 6.6.31.6 DS:
IBM.2107-1300861
ID
               FF
               1
Group
addrgrp
               Е
confgvols
               4
               0x1111
subsys
               3990-6
conbasetype
pprcconsistgrp Disabled
xtndlbztimout 120 secs
               300 secs
ccsesstimout
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. Can have one of the following values.

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 set path group ID fence state by the host.

Disabled

The volume is not set to the set path group ID fence state by the host.

N/A The volume is not capable of being set into the set path fence state by the host. For example, an alias volume.

CKD logical volume specific commands

Commands are referenced for tasks that are associated with System z count key data (CKD) logical volumes.

The following CKD logical volume specific commands are available:

chckdvo1

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 unit 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 z Systems CKD alias volumes (referred to as parallel access volumes or PAVs) in a storage unit.

mkckdvol

Creates z Systems 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 unit.

showckdvo1

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 *PG*0.

-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 (for a DS8000 model) Specifies the logical control unit number, 00 - FE
- *LL* (for a DS6000 model) Specifies the logical control unit number, 00 - 1F
- *VV* (for DS6000 and DS8000 models) 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. This command is not supported on DS6000 models.



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. This command is not supported on DS6000 models.

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.) The **-action** *releasespace* command cannot be used on a standard volume.

-quiet

(Optional) Turns off the confirmation prompt for the -action parameter.

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.

Isckdvol

The **lsckdvol** command displays a list of count key data (CKD) base and alias volumes in a storage image and status 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 **-1** parameters together.

-1

(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 **-1** 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 and 00 - 1F for the DS6000.

Note: Including the **-lcu** parameter can significantly reduce response time because the entire storage unit 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.
- **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. **Legacy** is always the reported value for a DS6000 model.

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 (for DS8000)

Specifies the logical control unit number, 00 - FE

LL (for DS6000)

Specifies the logical control unit number, 00 - 1F

VV (for both the DS8000 and DS6000)

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.

Note: If a column heading applies to the DS8000 only, a " - " value is displayed when the report is generated for a DS6000 model.

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

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 volume capacity (in cylinders).

Note: The deviceMTM 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 encryption group of an extent pool by using the **lsextpool -1**, 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. An inquiry on a DS6000 model always reports this value.

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).

Note: A value of 0 is displayed for the DS6000.

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. **Legacy** is always the reported value for a DS6000 model.

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.

manageckdvol

The manageckdvol command starts a process to initiate a change on count key data (CKD) volumes.





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

(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 "cancelled" 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, then all of the volumes that are still in the process of 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 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.

-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) Changes the extent pool ID of the volume so that the volume can migrate 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 can be available for I/O and copy services during migration. Its configstate can indicate that it is migrating.

-tier ssd ent n1 n1exclude

(Optional) Specifies which tier the volume is assigned to. This option is required with the **-action tierassign** parameter.

- **SSD** Solid state device tier
- **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

SSD or Enterprise tiers but not a Nearline tier.

-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.

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 represent the following values:

```
LL (for a DS8000 model)
```

Specifies the logical control unit number, 00 - FE

- LL (for a DS6000 model) Specifies the logical control unit number, 00 - 1F
- VV (for DS6000 and DS8000 models) 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 System 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 (for a DS8000 model)

Specifies the logical control unit number, 00 - FE

LL (for a DS6000 model)

Specifies the logical control unit number, 00 - 1F

VV (for DS6000 and DS8000 models)

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 System 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 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.

-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 zeroes, 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 represent the following values:

LL (for a DS8000 model)

Specifies the logical control unit number, 00 - FE

LL (for a DS6000 model)

Specifies the logical control unit number, 00 - 1F

VV (for DS6000 and DS8000 models)

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 neither the **-datatype** nor the **-3380** parameters are 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 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 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 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 wild cards:

- (#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 be 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 allocate 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: To use this subparameter, you must have previously created space-efficient storage (using the **mksestg** command) for the extentpool.

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 IBM System Storage DS8000 Thin Provisioning.

Note: To use this subparameter, you must have previously created space-efficient storage (using the **mksestg** command) for the extentpool.

-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 *PG*0.

-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 represent the following values:

- LL (for a DS8000 model) Specifies the logical control unit number, 00 - FE
- *LL* (for a DS6000 model) Specifies the logical control unit number, 00 - 1F
- VV (for DS6000 and DS8000 models)

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 (for a DS8000 model)

Specifies the logical control unit number, 00 - FE

LL (for a DS6000 model)

Specifies the logical control unit number, 00 - 1F

VV (for DS6000 and DS8000 models)

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.

For DS8000, 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 volume. This command can also be used to display the performance metrics for an individual volume ID.

► showckdvol-

-dev— *storage image ID*___ -metrics___

b						
-					volume_1D	
	- rank		L_tior_	└─ -volgrp── <i>volume group ID</i> ─		
	-Tulik	-pacingi p	- 61 61	-vorgip vorume_group_rb		

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 unit 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. There is only one volume group for CKD volumes 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 (for a DS8000 model)

Specifies the logical control unit number, 00 - FE

LL (for DS6000 model)

Specifies the logical control unit number, 00 - 1F

VV (for both the DS8000 and DS6000 model) 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)
TSE	1.7	-	3339	0.0

realextents	virtualextents	migrating	migratingfrom	perfgrp	resgrp
1	0	0	-	PG0	RG0

tierassignstatus	tierassignerror	tierassignorder	tierassigntarget	%tierassigned
Assigning	-	ETdata	SSD	54

-----Rank extents------

Rank	Extents
R0	1
R2	2

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 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.

Virtual space fault

Indicates that the logical volume has a space-efficient storage allocation method, and there was not enough space 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 designations can be 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 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 encryption group of an extent pool by using the **lsextpool -1**, 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 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. An inquiry on a DS6000 model always reports this value.

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.

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) and Y represents tenths of a GiB, which is limited to a single digit (0 - 9).

Note:

- 1. A " " value is displayed in this column if the value displayed in the SAM column is Standard or ESE.
- 2. A " " value is displayed for the DS6000.

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. Legacy is always the reported value for a DS6000 model.

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 " - " 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).

Note: A value of 0 is displayed for the DS6000.

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.

migrating

The number of extents for this volume that are currently being migrated.

migratingfrom

A list of one or more extent pool IDs where the extents are migrating from. If there are no migrating extents, 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.

- SSD Solid-state device tier
- **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.

NLExcluded

SSD or 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:

уууу	the year
MM	the month (01-12)
dd	the day (01-31)
Т	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.

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.

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 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 issued 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

segreadhits

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

NVSspade1

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	myvo10800
ID	0800
accstate	Online
datastate	Normal
configstate	Normal
•••	
migrating	20
perfgrp	PGO
migratingfrom	P0
resgrp	RG1
tierassignstatus	Assigning
tierassignerror	-
tierassignorder	ETdata
tierassigntarget	SSD
%tierassigned	54
etmonpauseremain	1H44M
etmonitorreset	2013-07-26T14:00:00+07
Ti	ier Distribution ==================
Tier %allocated	
SSD 54	
ENT 46	

Report field definitions (*-tier* **parameter is specified)**

Tier ID

SSD Solid-State Device tier.

- **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.

NLExcluded

SSD or Enterprise tiers but not a Nearline tier.

Unknown

Tier is unknown.

%allocated

Percentage of volume capacity on this tier.

Example

Specifying the -pathgrp parameter

If you specify the **-pathgrp** parameter and there are no path groups for this volume, the following message is displayed.

CMUC00234I lsckdvol: No Path Groups found.

If you specify the **-pathgrp** parameter and there are path groups for this volume, a path group status table is appended to the resulting output.

dscli> showckdvol -pathgrp efff

The resulting output

Name ID accstate datastate configstate	efff EFFF Online Normal Normal				
migrating perfgrp migratingfrom resgrp	0 PG31 - RG62	Dath Grou	10 status		
GroupID		State	Reserve	Mode	Sysplex
800002B9472827CA7 880005B9472827CAA 800009B9472827CAC	D6FBA	Ungrouped Grouped Ungrouped	Disabled Disabled Disabled Disabled	Single Multi-path Single	====== N/A N/A 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 breakdown the ID into distinct fields with the following partial display of the ICKDSF logical path status table.

+	PATH GI	ROUP ID
ID	SERIAL	CPU CPU TIME TYPE STAMP
800002	B947	2827 CA78BC17
880005	B947	2827 CAAD6FBA
800009		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.


1s1ss Generates a report that displays a list of logical subsystems (LSSs) for a storage unit and the status information for each logical subsystem in the list.

show1ss

Generates a report that displays the detailed properties of a specified LSS.

chiss

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.

-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 for the DS8000 and it is 00 - 1F for the DS6000.

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.2017-75FA120 -extlongbusy 120 06 0F
```

The resulting output

```
LSS 06 successfully modified.
LSS 0F successfully modified.
```

Islss

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 -1 and the -s parameters together.

-1

(Optional) Displays the default output plus the SSID and resource group ID. You cannot use the **-1** 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 for the DS8000 and 00 - 1F for the DS6000.

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.1750-75FA120 -1

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

\mathbf{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 for the DS8000 and 00 - 1F for the DS6000.

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 **-1** 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
IBM.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 for the DS8000 and 00 - 1F for the DS6000.

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: March 28, 2012 3:18:07 PM MST IBM DSCLI
Version: 7.6.20.138 DS: IBM.2107-75LH321
ΙD
               14
               0
Group
               1
addrgrp
stgtype
               fb
confgvols
               16
subsys
               0xFF14
pprcconsistgrp Disabled
xtndlbztimout 60 secs
resgrp
               RG0
======Soft Fence State========
         ID sfstate
Name
------
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
RegrFBVoll 140A Enabled
```

RegrFBVoll 140B Enabled RegrFBVoll 140C Enabled RegrFBVoll 140D Enabled RegrFBVoll 140E Enabled RegrFBVoll 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 unit 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 unit.

rmfbvol


Deletes one or more specified fixed block volumes from a storage unit.

showfbvo1

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

-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.

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 (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 *PG*0.

-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 1.
- *XY* This parameter is not available for DS6000 models. Specifies the logical subsystem number, 00 FE.
- *XY* This parameter is only available for DS6000 models. Specifies the logical subsystem number, 00 1E.
- 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 -type ess -cap 8.6 000a

The resulting output

CMUC00026I chfbvol: FB volume 000A 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. This command is not supported on DS6000 models.

► — initfbvol-	
▶ volume_IL	> 4

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.) The **-action** *releasespace* command cannot be used on a standard volume.

-quiet

(Optional) Turns off the modification confirmation prompt for the **-action** parameter.

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 volumes in a storage image and status information for each volume in the list.



Parameters

Note: For a storage unit 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, 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 IDs. You cannot use the -1 and the -s parameters together.

-1

(Optional) Displays default output plus additional attributes that are identified as long output in the Report field definitions list. You cannot use the **-1** 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 zeroes, 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 for the DS8000 and 00 - 1F for the DS6000.

-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.
- 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. 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

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 4 hexadecimal digits in the form of *XYZZ* where:

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1, for DS6000 and 0-F for DS8000.

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
My_ volume_ 0002	0102	Online	Normal	Normal	2107-A07	FB 520P
My_ _volume 0003	0103	Online	Normal	Normal	2107-900	FB 512

Name	ID	accstate	datastate	config state	device MTM	data type
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

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. There was not enough space 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 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 has been 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 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 has been lost. As a result, some of the data 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 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

Indicates the volume device type and model. The volume MTM (machine type, model) is determined by the fixed block volume data type and the volume 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 encryption group of an extent pool by using the **lsextpool -1**, 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. An inquiry on a DS6000 model always reports a value of standard.

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 there are no volume groups that are associated with the volume.

Unknown is displayed if information about the volume groups is not available.

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. **Legacy** is always the reported value for a DS6000 model.

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).

Note: A value of 0 is displayed for the DS6000.

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.

managefbvol

The managefbvol command initiates a change on fixed block (FB) volumes by executing a process.



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

(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 "cancelled" 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 "cancelled" 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.

-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 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.

-tier ssd|ent|nl |nlexclude

(Optional) Specifies which tier the volume is assigned to. This option is required with the **-action tierassign** parameter.

SSD Solid state device tier,

- **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

SSD or Enterprise tiers but not a Nearline tier.

-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.

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 represent the following values:

LL (for a DS8000 model)

Specifies the logical control unit number, 00 - FE

LL (for a DS6000 model)

Specifies the logical control unit number, 00 - 1F

VV (for DS6000 and DS8000 models)

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 managefbvol command


dscli> managefbvol -dev IBM.2107-75FA120 -action migstart -extpool P2 0100

The resulting output

CMUC00000I managekbvol: FB Volume 0100 action migstart executed successfully.

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 have not yet been created, create temporary volume groups and assign new fixed block volumes to the temporary 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 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 system. Using the **-dev** parameter will temporarily override any defined value for *devid* for the current command.

-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.

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

The storage sizes and the data types for this volume are listed in the following table:

Decimal gigabyte (GB) is 10⁹ 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.

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 DS8000 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 1 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 one 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 issue 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.

Note: To use this subparameter, you must have previously created space-efficient storage (using the **mksestg** command) for the extentpool.

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 IBM System Storage DS8000 Thin Provisioning.

Note: To use this subparameter, you must have previously created space-efficient storage (using the **mksestg** command) for the extentpool.

-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 *PG*0.

-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 using this command, the LSSs are assigned to the same resource group as the logical volumes.

-t10dif

(Optional) Specifies that the DS8000 system is using the CRC-16-T10-DIF algorithm 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 represent the following values:

X (for DS8000 and DS6000)

Specifies the address group, 0 -1 for DS6000 and 0-F for DS8000.

XY (DS6000 only)

Specifies the logical subsystem number, 00 - 1E.

XY (DS8000 only)

Specifies the logical subsystem number, 00 - FE.

ZZ (for DS8000 and DS6000)

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:

For DS8000:

dscli> rmfbvol -dev IBM.2107-75FA120 0100-010F 0180-018F 0120

For DS6000:

dscli> rmfbvol -dev IBM.1750-13ABR4A 0005-00FF

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:

For DS8000:

dscli> rmfbvol IBM.2107-75FA120/0100-010F 0180-018F 0120

For DS6000:

dscli> rmfbvol IBM.1750-13ABR4A/0005-00FF

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

X (for DS6000 and DS8000 models)

Specifies the address group, 0 -1 for DS6000 and 0-F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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 0009 successfully deleted.

CMUC00028I rmfbvol: FB volume 0009 successfully deleted.


CMUC00028I rmfbvol: FB volume 0008 successfully deleted.

CMUC00028I rmfbvol: FB volume 0008 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 0011 successfully deleted.
CMUC00028I rmfbvol: FB volume 0012 successfully deleted.
CMUC00028I rmfbvol: FB volume 0013 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 volume. This command can also be used to display the performance metrics of a fixed block 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 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 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, which lists the set of tiers that have storage allocated 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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0–F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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 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 showfbvol 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 6000

The resulting output

Name	ID	acc state	data state	config state	device MTM	data type	addrgrp
My_ volume _6000	6000	Online	Normal	Normal	2107-900	FB 512	6

extpool	exts	captype	cap (2^30B)	cap (10^9B)	cap (blocks)	volgrp	ranks	dbexts
P0	1622	DS	1622.0	-	3401580544	-	2	0

sam	repcapalloc	eam	reqcap (blocks)
Standard	-	rotateexts	2097152

realextents	virtualextents	migrating	migratingfrom	perfgrp	resgrp
1	0	0	-	PG0	RG0

tierassignstatus	tierassignerror	tierassigntarget	%tierassigned
Assigning	-	SSD	54

-----Rank extents-----

Rank	Extents
R0	1
R2	2

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. There was not enough available space 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 has occurred:

Data states that do not apply:

- Rank failed
- Rank repairing
- Rank repaired
- Global inaccessible
- · Global lost data

Conditions - one of the following conditions 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 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 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 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 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 de-allocating 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 "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 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 re-drive 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 volume 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 encryption group of an extent pool by using the **lsextpool -1**, or **showextpool** commands.

Exts

Indicates the number of real and virtual extents used by the designated volume ID.

Captype

Indicates capacity unit type 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 iSeries.

Cap (2^30B)

Indicates the size of the volume that is available for host system access in gigibytes (GiB).

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 gigabytes (GB).

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 there are no volume groups that are associated with the volume.

Unknown is displayed 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. An inquiry on a DS6000 model always reports this value.

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 deallocated from the extent pool.

Note: IBM Database protection feature supports standard volumes only.

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:

1. A " – " value is displayed in this column if the value displayed in the SAM column is not TSE.

2. A " – " value is displayed for the DS6000.

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. **Legacy** is always the reported value for a DS6000 model.

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).

Note: The value 0 is displayed for the DS6000.

realextents

Indicates the number of real extents used by the logical volume.

virtualextents

Indicates the number of virtual extents used by the logical volume.

migrating

The number of extents for this volume that are currently being migrated.

migratingfrom

A list of one or more extent pool IDs where the extents are migrating from. If there are no migrating extents, 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 has 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 see if 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 have 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

Method 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.

SSD Solid state device tier

ENT Enterprise tier consists of drives with speeds of 10K RPM, 15K RPM, or a combination of both.

NL Near Line tier consists of high volume SATA or SAS Near Line disk drives.

NLExcluded

Any tier except NL tier.

Unknown

Assign action was specified, but the target tier is unknown.

" – " No assign action was specified.

%tierassigned

The percentage of the volume 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:

уууу	the year
MM	the month (01-12)
dd	the day (01-31)
Т	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.

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.

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.

NVSspade1

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.

mulcate

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 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

Invoking the showfbvol to show tier statistics

dscli> showfbvol -tier 0000

The resulting output

Namemyvol0800ID0800accstateOnlinedatastateNormalconfigstateNormal...migrating20

perfgrp	PGO
migratingfrom	P0
resgrp	RG1
tierassignstatus	Assigning
tierassignerror	-
tierassignorder	ETdata
tierassigntarget	SSD
%tierassigned	54
etmonpauseremain	1H44M
etmonitorreset	2013-07-26T14:00:00+07
•	ier Distribution ==============
Tier %allocated	
SSD 54	
ENT 46	

Report field definitions (-tier parameter is specified)

- Tier ID
 - **SSD** Solid State Device tier.
 - **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.

NLExcluded

SSD or Enterprise tiers but not a Nearline tier.

Unknown

Tier is unknown

%allocated

The percentage of volume capacity on this tier.

Example 4

Specifying the -pathgrp parameter

If you specify the -pathgrp parameter and there are no path groups for this volume, the following message is displayed.

CMUC00234I lsfbvol: No Path Groups found.

If you specify the **-pathgrp** parameter and there are path groups for this volume, a path group status table is appended to the resulting output.

dscli> showfbvol -pathgrp efff

The resulting output

Name ID accstate datastate configstate	efff EFFF Online Normal Normal				
 migrating perfgrp migratingfrom	0 PG31				
resgrp	- RG62				
	======P	ath Group	status====		
GroupID	St	ate	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 eleven-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 the hosts are free to define, or re-define, the meaning of this value with no impact to the DS8000 system. However, some programs such as ICKDSF break down the ID into distinct fields with the following partial display of the ICKDSF logical path status table.

++-	
PATH GROUP ID	•••
ID SERIAL TYPE STAMP	•••
800002 B947 2827 CA78BC17	•••
880005 B947 2827 CAAD6FBA	••••
800009B947 2827CAC684B9	••••

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 there are SCSI reservations 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 PG0 perfgrp migratingfrom RG0 resgrp ======SCSI Reserve Status====== PortID WWPN ReserveType

======		
10040	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 World Wide Port Name displayed as sixteen 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 unit and the status information for each volume group in the list.

mkvolgrp

Creates a volume group in a storage unit.

rmvolgrp


Deletes the specified volume group or volume groups from a storage unit.

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.

For DS8000, 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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0-F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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.
- 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
```

If the **-action** *replace* parameter is specified without specifying the **-lun** parameter, lun_ID=00 is assumed.

Volume_Group_ID | -

0001 : 11

(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.

For DS8000, example: IBM.2107-75FA120

- S

(Optional). Displays volume group IDs only. You cannot use the -1 and the -s parameters together.

-1

(Optional). Displays the default output. You cannot use the **-1** 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 4 hexadecimal digits in the form of *XYZZ* where:

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0-F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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 -1
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Device: IBM.2107-75FA120

Name	ID	Туре
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.

 \mathbf{ID}^{\star}

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.

Туре

Indicates the type of the volume group.

Key:

- * Displayed when the **-s** parameter is specified.
- + Displayed only when the -1 parameter is specified.

mkvolgrp

The **mkvolgrp** command creates a volume group in a storage image.



		Volume Group Name
└─ -volume── <i>volume_ID</i> ── · · · · ┘	└─ -lun── <i>lun_ID</i> ─┘	

Parameters

Notes:

- 1. The DS8000 system assigns the ID 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.
- **3**. When you create DS6000 volume groups for (RedHat) Linux using the **mkvolgrp** command, the **-type** parameter must be set to *scsimap256*.
- 4. When you create DS6000 volume groups for AIX5L, the **-type** parameter must be set to *scsimap*.

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.

For DS8000, 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 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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0-F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

- *XY* (for a DS6000 model) Specifies the logical subsystem number, 00 - 1E.
- ZZ (for DS6000 and DS8000 models)

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.

For DS8000, 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: For DS8000, 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: For DS8000,

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

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.

For DS8000, 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.

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:

For DS8000: 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:

For DS8000:

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	Туре	Vols
My_host_ system_ volumes		SCSI- MAP 256	1000 1001 1002 1003 1004 1005
Volumes		200	1004 1005 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	Туре	Vols
myVG1	V18	SCSI-	1000 1001
		Mask	1002 1003
			1004 1005
			1006 1007

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*.

Туре

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. This command is not supported on DS6000 models.

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 4 hexadecimal digits in the form of *XYZZ* where:

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0-F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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.

For DS8000, example: IBM.2107-75FA120/0001

Example

Invoking the clearvol command


dscli> clearvol -dev IBM.2107-75FA120 0001

The resulting output

Volume 0001 successfully cleared.

Isvolinit

The **lsvolinit** command displays volumes that are being initialized using flash init. CKD volumes are not displayed. This command is not supported on DS6000 models.



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 **-1** parameters together.

-1

(Optional). Displays the default output. You cannot use the **-1** 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 **-1** parameter is specified.

Space-efficient storage commands

Space-efficient 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.

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.

▶▶-	- chsestg dev— storage_image_ID captype— gb _ blocks- _ cyl _ mod1
▶	reptype─ tse─repcap─_ <i>capacity</i> reppercent─_ <i>percentage</i>
•	repcapthreshold_ repository_threshold_percentagequietwait
▶─	- extentpool_ID

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 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).

Select one of the following capacity unit types:

gb (Gibibyte)

Specifies the capacity unit type as gibibytes (GiB). 1 GiB = 2 097 152 blocks.

- **block** Specifies the capacity unit type as blocks. This capacity unit type can be selected only for fixed block storage. 1 block = 512 Bytes.
- **cyl** Specifies the capacity unit type as cylinders. This capacity unit type can be selected only for CKD storage. 1 cylinder = 15 tracks.
- **Modl** Specifies the capacity unit type in Mod1 units. This capacity unit type can be selected only for CKD storage. 1 Mod1 = 1113 cylinders.

-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.

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 be at least 2:1.
- 2. Only one of the following parameters can be specified: -repcapthreshold, -repcap, -reppercent, and -vircap.

-reptype tse | ese

(Optional) Specifies the type of repository that is affected by the **repcap**, **reppercent**, and **repcapthreshold** parameters. If the type is not specified, the **-reptype** parameter defaults to tse.

Valid repository types are as follows:

tse

Specifies a track space-efficient physical repository. This type is the default.

Note: A TSE repository is required to create TSE volumes.

ese

Specifies an extent space-efficient physical repository.

Note: An ESE repository is not required to create ESE volumes.

-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 TSE repositories:

- A TSE repository is required to create TSE volumes.
- Changing the capacity size of an existing TSE repository is not supported.
- The **chsestg** command cannot be used to create a TSE repository where one does not exist. Use the **mksestg** command.
- The **chsestg** command can be used to remove a TSE repository. Use either the **-repcap** or the **-reppercent** parameters to remove the existing repository by specifying 0 (zero) as the parameter value. However, using the **rmsestg** command to remove a TSE repository is recommended.
- All TSE volumes must be removed before attempting to remove the TSE repository.

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, and 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.
- 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 repository can be resized to any capacity from the size of the capacity used to provision any existing ESE volumes to just 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, 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 and 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 set manually to a higher value, the total virtual capacity is usually equal to the sum of all the capacities of the 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 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.

-quiet

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

-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 only 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.
```

Issestg

The lssestg command displays a list of the space-efficient repositories in the storage unit.



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 **-1** parameters together.

-1

(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 **-1** 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 unit.

dscli> lssestg -dev IBM.2107-75FA120 -l

The resulting output

```
dscli> lssestg -l
Date/Time: September 24, 2013 9:21:08 AM MST IBM DSCLI
Version: 7.7.20.524 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	tste	0.0	0.0	0.0
1040.4	1600.0	ese	0.0	500.0	0.5

Notes:

- 1. The extent pool P5 has a TSE repository, but no ESE repository, and an over-provisioned ratio of 1.5 to 1.
- 2. 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.
- **3**. The extent pool P13 has the virtual capacity that is specified but neither a TSE or an ESE repository. ESE volume creation is still allowed. However, because there is no ESE repository, vircapalloc is always zero. See the virconfigured field of the **showextpool** command to see total space-efficient storage of the extent pool.
- 4. 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 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 there are internal inconsistencies 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 there are 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 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 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 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.

bel ow

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 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 **-1** 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 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.

-captype gb | blocks | cyl | mod1

(Optional) Specifies the unit type of capacity 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.
- **block** 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.

-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. If the type is not specified, the **-reptype** parameter defaults to tse.

The following repository types are valid:

tse

Specifies a track space-efficient physical repository. This type is the default.

Note: A TSE repository is required to create TSE volumes.

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 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 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 currently 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 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 already 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.

An error occurs if you use this command to create a TSE/ESE repository when a TSE/ESE repository already exists.

For TSE repositories:

- A TSE repository is required to create TSE volumes.
- All storage 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, which is equivalent to 16 GiB.
 - **cyl** = 16 740 cylinders.

For ESE repositories:

- An ESE repository is not required to create ESE volumes.
- All storage 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 will become a part of the newly created ESE repository but will remain otherwise unchanged. Any existing ESE volumes will 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, specified as a percentage of the total virtual capacity. Creating a repository that already 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 already exists.

This parameter is provided as a convenience and 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 usually equal to the sum of all the capacities of the space-efficient volumes.

Note:

1. 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 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

Invoking the mksestg command to create space-efficient storage in an extent pool.


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.

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:

a11

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.
- 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.

Valid repository types are as follows:

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	Р3
stgtype	fb
datastate	Normal
	Normal
configstate	below
repcapstatus	
%repcapthreshold	0
repcap(GiB)	100.0
repcap(Mod1)	-
repcap(blocks)	209715200
repcap(cyl)	-
<pre>repcapalloc(GiB/Mod1)</pre>	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
regrepcap(GiB/Mod1)	100.0
reqvircap(GiB/Mod1)	1000.0
reptype	ese
opratio	0.5
0010010	

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

Indicates the type of repository, as follows:

tse

Indicates a track space-efficient physical repository.

ese

Indicates an extent space-efficient physical repository.

- 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).

I/O Priority Management commands

The performance group attribute on logical volumes associates those volumes with a performance group object. Each performance group has an associated performance policy that determines how the I/O priority manager processes I/O operations for the logical volume.

The I/O priority manager maintains statistics for the set of logical volumes in each performance group that can be queried using the **lsperfgrprpt** command. If management is performed for the performance policy, the I/O priority manages the I/O operations of all managed performance groups to achieve the goals of the associated performance policies. If not specified, the performance group defaults to 0. The following table provides the performance groups that are predefined and have the associated performance policies:

Performance group	Performance policy	Performance policy description
0	0	No management
1-5	1	Fixed block high priority
6-10	2	Fixed block medium priority
11-15	3	Fixed block low priority
16-18	0	No management
19	19	CKD high priority 1
20	20	CKD high priority 2
21	21	CKD high priority 3
22	22	CKD medium priority 1
23	23	CKD medium priority 2
24	24	CKD medium priority 3
25	25	CKD medium priority 4
26	26	CKD low priority 1
27	27	CKD low priority 2
28	28	CKD low priority 3
29	29	CKD low priority 4
30	30	CKD low priority 5
31	31	CKD low priority 6

This section contains commands that are used to manage quality of service for DS8000 models only.

The following I/O Priority Management commands are available:

lsperfgrp

Displays a list of performance groups and information for each performance group in the list.

lsperfgrprpt

Displays a list of performance group statistics.

lsperfrescrpt

Displays a list of performance resources and information for each performance resource in the list.

Isperfgrp

The **lsperfgrp** command displays a list of performance groups and information for each performance group in the list.

►►—lsperfgrp-

___performance_group_ID___

Parameters

-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, 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 performance group ID. You cannot use the **-s** and the **-1** parameters together.

-1

(Optional) Displays the default output. You cannot use the **-s** and the **-1** parameters together.

-pol policy_ID

(Optional) Displays only the performance groups with the specified policy ID.

performance_group_ID ... | -

(Optional) Displays only the performance groups with the performance group IDs specified. Multiple IDs or ID ranges must be separated with a white 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.

Example

Invoking the lsperfgrp command

dscli> lsperfgrp

The resulting output

ID	pol
PG0	1
PG1	2
PG2	2

Report field definitions

\mathbf{ID}^{*}

Displays the unique identifier that is assigned to this performance group ID.

po1

Displays the policy ID number.


Key:

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

+ Displayed only when the **-1** parameter is specified.

Isperfgrprpt

The **lsperfgrprpt** command displays a list of performance reports for the specific set of performance groups, or all if none are specified.



Parameters

-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) Use this parameter to display only the performance group IDs. You cannot use the **-s** and the **-1** parameters together.

-1

(Optional) Use this parameter to display the default output and extra attributes that are identified as long output. You cannot use the **-s** and the **-1** parameters together.

-dapair dapair_ID

(Optional) Displays only the performance groups with the specified DA pair ID. You cannot use the **-dapair** and the **-rank** parameters together.

-rank rank_ID

(Optional) Displays only the performance groups with the specified rank ID. You cannot use the **-dapair** and the **-rank** parameters together.

-start time

(Optional) Specifies the start time of the report in the past, relative to the current time. The time format is specified in days, hours, minutes; for example: 1d,2h,3m. The default is 1h, which means 1 hour before the current time.

-stop time

(Optional) Specifies the stop time of the report in the past, relative to the current time. The time format is specified in days, hours, minutes. The default is 0m, which means the current time.

-interval time

(Optional) Specifies the interval of time between report samples. The time format is specified in days, hours, minutes. The default is 5m, which means a 5-minute interval between samples.

performance_group_ID | -

(Optional) Displays only the performance groups with the performance group IDs specified. Multiple IDs or ID ranges must be separated with a white space between each value. If you use the dash (-), the specified value is read from standard input.

Note: Each line of the display is defined as a "report" for that resource at that time. This command has a maximum limit of 256 reports, but is configurable in the dscli.profile. If the limit is exceeded,

the command displays an error message and the limit number of reports. The truncated reports are not always at the end of the displayed reports, but might appear as missing reports within the range of the displayed reports.

Example

Invoking the lsperfgrprpt command

dscli> lsperfgrprpt

The resulting output

Date/Time	grp	resrc	avIO	avMB	avresp	pri
2009-05-20/ 15:23:19	PG50	IBM.2107- 75FA120	123465	1.234	12.345	12

aveQ	tgtQ	%hlpT	%dlyT	%impT	mnIO	mxIO	%idle	mnMB	mxMB
123	123	100	100	100	123456	123456	100	1.234	1.234

mnresp	mxresp	%Hutl	%VHutl	%loQ	%hiQ	%hlp#	%rep	%dly#	%ceil
123.450	123.450	100	100	100	100	100	100	100	1234

Report field definitions

Date/Time

Indicates the time stamp for the performance group statistics.

- grp Indicates the performance group ID.
- **resrc** Indicates the resource ID for the specified resource.
- **avIO** Indicates the average (mean) I/O operation per second.
- avMB Indicates the average (mean) megabytes per second transferred.
- **avresp** Indicates the average (mean) response time in milliseconds (ms) for track I/O operations during this interval.
- pri Indicates the performance group priority.
- **avQ** Indicates the average (mean) I/O-weighted QoS index during this interval.
- tgtQ Indicates the QoS target value for the performance group.

%hlpT

Indicates the percent of intervals in which I/Os were helped.

%dlyT

Indicates the percent of time in which I/Os were delayed.

%impt

Indicates the percentage of impact on those I/Os delayed.

- mnIO Indicates the minimum track I/O operation per second.
- mxIO Indicates the maximum track I/O operation per second.
- %idle Indicates the percentage idle.

mnMB

Indicates the minimum megabytes per second transferred.

Indicates the maximum megabytes per second transferred.

mnresp

Indicates the minimum response time for track I/O operations during this interval.

mxresp

Indicates the maximum response time for track I/O operations during this interval.

%Hutl Indicates the percentage of time in which the specified resource had utilization high enough to warrant workload control.

%VHutl

Indicates the percentage of time in which the specified resource had very high utilization.

- %loQ Indicates the interval percentage with low QoS.
- %hiQ Indicates the time percentage with high QoS.
- %hlp# Indicates the percent of I/Os which were helped.


%req Indicates the percentage of time in which I/O help was requested.

%dly# Indicates the percentage of I/Os delayed by throttling.

%ceil Indicates the ceiling for performance group on acceptable impact.

Isperfrescrpt

The **lsperfrescrpt** command displays a list of performance reports for a given resource or set of resources of a given type.



Parameters

-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, 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) Use this parameter to display only the performance group IDs. You cannot use the **-s** and the **-1** parameters together.

-1

(Optional) Use this parameter to display the default output and additional attributes that are identified as long output. You cannot use the **-s** and the **-1** parameters together.

-start time

(Optional) Specifies the start time of the report in the past, relative to the current time. The time format is specified in days, hours, minutes; for example: 1d,2h,3m. The default is 1h, which means one hour before the current time.

-stop time

(Optional) Specifies the stop time of the report in the past, relative to the current time. The time format is specified in days, hours, minutes. The default is 0m, which means the current time.

-interval time

(Optional) Specifies the interval of time between report samples. The time format is specified in days, hours, minutes. The default is 5m, which means a five-minutes interval between samples.

dapair_id | rank_id | -

(Optional) The specified resource for which performance reports should be displayed.

A device adapter pair ID (dapair_id) is a decimal number prefixed by the letters DP. A rank number (rank_id) is a decimal number prefixed by the letter R.

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. If specifying a resource, only one resource may be specified.

Note: Each line of the display is defined as a "report" for that resource at that time. This command has a maximum limit of 256 reports, but is configurable in the dscli.profile. If the limit is exceeded, the command displays an error message and the limit number of reports. The truncated reports are not always at the end of the displayed reports, but might appear as missing reports within the range of the displayed reports.

Example

Invoking the lsperfrescrpt command

dscli> lsperfrescrpt DP2

The resulting output

Date/Time	resrc	avIO	avMB	avresp	%Hutl	%hlpT	%dlyT	%impT
2009-05-20/15:23:19	IBM.2107- 75FA120	555	2.046	1.477	0	13	0	212

Report field definitions

Date/Time

Indicates the time stamp for the performance group statistics.

- resrc Indicates the resource ID for the specified resource.
- **avIO** Indicates the average (mean) IO operation per second.
- **avMB** Indicates the average (mean) megabytes per second transferred.
- **avresp** Indicates the average (mean) response time in tenths of a second for track IO operations during this interval.
- **%Hutl** Indicates the percentage of time in which the specified resource had utilization high enough to warrant workload control.

%hlpT

Indicates the percent of intervals in which IOs were helped.

%dlyT

Indicates the percent of time in which IOs were delayed.

%impt

Indicates the percentage of impact on those I/Os delayed.

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.

►►(chresgrp—	-dev- s	torage_image_	ID	r	esource_group_	label	
▶	-name— r	resource_grou	up_name_	resou 	rce_group_ID-]		 _►

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.

Isresgrp

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 **-1** and the **-s** parameters together.

-1

(Optional) Displays default output. You cannot use the -1 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

\mathbf{ID}^{\star}

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.

Labe1

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 -1 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 encryption 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 resource 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 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.

▶ — rmresgrp –

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.

►→ showresgrp _____ resource_group_ID ____ resource_group_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 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.

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.

▶▶—	commitflash	→
▶	source_volume_ID	≯ ∢

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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0-F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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.

For DS8000, 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 IBM Remote Pair Copy option that you want to use. The IBM 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 IBM Remote Pair Copy option for FlashCopy operations when possible. The IBM 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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0-F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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.

For DS8000, 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> 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 -1 parameters together.

-1

(Optional) Displays the default output plus additional attributes that are identified as long output. You cannot use the **-s** and the **-1** 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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0–F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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 -1 0100:0200 0101:0201 0102:0202 0103:0203

The resulting output

ID	SrcLSS	Sequen- ceNum	Time- out	Active Copy	Record- ing	Persis- tent	Revert- ible
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	Back- ground- Copy	OutOf- Sync- Tracks	Date- Created	Date- Synced	State	is- TgtSE	Pmir
Enabled	Disabled	Disabled	0	10/01 /2007 02:20:00	10/01 /2007 02:23:47	Valid	TSE	No
Enabled	Disabled	Disabled	0	10/01 /2007 02:20:00	10/01 /2007 02:23:47	Valid	TSE	No
Enabled	Disabled	Disabled	0	10/01 /2007 02:20:00	10/01 /2007 02:23:47	Valid	ESE	No
Enabled	Disabled	Disabled	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 16. 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. No value is displayed in the DS6000 models.

TargetWriteEnabled

Indicates whether this FlashCopy relationship was established with the allow target writes option. No value is displayed in the DS6000 models.

BackgroundCopy

Indicates whether this FlashCopy relationship was established with the run background copy option. No value is displayed in the DS6000 models.

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. No value is displayed in the DS6000 models.

DateCreated

Indicates the date and the time that the FlashCopy relationship was established. No value is displayed in the DS6000 models.

DateSynced

Indicates the date and time that FlashCopy relationship was synchronized, or specifies " - " if the relationship is not synchronized. No value is displayed in the DS6000 models.

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 IBM System Storage DS8000 Thin Provisioning.

Unknown

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

Pmir 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 unit. 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 can 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 can fail if the IBM 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 an 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.

mkflash

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

▶ mkflash	>
	F
-freeze -record -persist -multinc -tgtse -cp -wait	
└ -seqnum— <i>flash_sequence_num</i> _ └ -pmir— no— · resetreserve_ · required— · resetreserve_ · preferred	
- required-	
└─ preferred-J	
· ·	
▶ source_volume_ID:target_volume_ID	

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) 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 IBM Remote Pair Copy option that you want to use. The IBM 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 unit. 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 IBM Remote Pair Copy option for FlashCopy operations when possible. The IBM 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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0 - F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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.

For DS8000, 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) 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 IBM Remote Pair Copy option that you want to use. The IBM 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 unit. 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 IBM Remote Pair Copy option for FlashCopy operations when possible. The IBM 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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0 - F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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.

For DS8000, 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.

DS8000 example: IBM.2107-75FA120

DS6000 example: IBM.1750-685FA120

-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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0 - F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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.

>> —	rmflasl	n	- storage_image_ID_			-tgtreleasespace	
▶──	-cp		-resettgtinhibi	it	seqnum	– flash_sequence_number	>

```
SourceVolumeID:TargetVolumeID
```

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.
- 5. The **-seqnum** parameter is not supported for a 2105 machine type.

-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.

DS8000 example: IBM.2107-75FA120

DS6000 example: IBM.1750-68FA120

-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 z System host are removed from the specified target volume.

-tgtreleasespace

(Optional - DS8000 only) Specifies that you want the system to release the space that has been allocated 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

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

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 (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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 (for DS8000): 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 DS8000 value for the XX is 00 - FE and the DS6000 value is 00 - 1F.

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.

DS8000 example of a fully qualified LSS ID: IBM.2107-75FA120/00

DS6000 example of a fully qualified LSS ID: IBM.1750-68FA120/00

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.

>>	setflashrevertible—		
▶	-seqnum— flash_sequ	uence_num source_volume_ID:target_volume_ID	 ••

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.

For DS8000, example: IBM.2107-75FA120

For DS6000, example: IBM.1750-68FA120

-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).

This parameter is not supported for machine type 2105.

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 4 hexadecimal digits in the form of *XYZZ*, where:

```
XY (for a DS8000 model)
Specifies the logical subsystem number, 00 - FE.
```

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models) Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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.

For DS8000, 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 DS8000 value for the XX is 00 - FE and the DS6000 value is 00 - 1F.

-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 (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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.

For DS8000, 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.

-conduit— LSS_ID-

►►— resyncremoteflash-

└──-dev── storage image ID─┘

▶-		→
▶-	-tgtse -cp -seqnum flash_sequence_num -srcss SS_ID	→
▶	-resetreserve SourceVolumeID:TargetVolumeID	→4

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.

For DS8000, 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, for the DS8000, and 00 - 1F, for the DS6000.

For DS8000, 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 (for a DS8000 model) Specifies the logical subsystem number, 00 - FE.
- XY (for a DS6000 model) Specifies the logical subsystem number, 00 - 1E.
- ZZ (for DS6000 and DS8000 models) Specifies the volume number, 00 - FF.
- X (for DS6000 and DS8000 models) Specifies the address group, 0 - 1.

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.

For DS8000, 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.

Isremoteflash

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 (issue lsremoteflash 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 (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

- S

(Optional) Displays only FlashCopy pair IDs. You cannot use the -1 and the -s parameters together.

-1

(Optional) Displays the default output plus out-of-sync tracks and date that the FlashCopy relationship was created. You cannot use the **-1** 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. This parameter is not supported on DS6000 models.

-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 (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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. No value is displayed for the DS6000.

TargetWriteEnabled

Indicates (enabled or disabled) whether this FlashCopy relationship was established with the Allow Target Writes option. No value is displayed for the DS6000.

BackgroundCopy

Indicates (enabled or disabled) whether this FlashCopy relationship was established with the Run Background Copy option. No value is displayed for the DS6000.

OutofSyncTracks⁺

Indicates the number of tracks that are not synchronized for this FlashCopy relationship. No value is displayed for the DS6000. A dash (-) is displayed when the track counter is not available.

DateCreated⁺

Indicates the date and the time that the FlashCopy relationship was established. No value is displayed for the DS6000.

DateSynced

Indicates the date and the time that this FlashCopy relationship was synchronized, or " - " if the relationship is not synchronized. No value is displayed for the DS6000.

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.

Note: No value is displayed for the DS6000.

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 IBM System Storage DS8000 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 unit. 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 an 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 **-1** 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 DS8000 value for the XX is 00 - FE and the DS6000 value is 00 - 1F.

For DS8000, 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 (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models) Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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.

For the DS8000, 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 full qualified LSS ID. The fully qualified LSS ID format is storage_image_ID/XX. The DS8000 value for the XX is 00 - FE and the DS6000 value is 00 - 1F.

For DS8000, 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 (for a DS8000 model)
Specifies the logical subsystem number, 00 - FE.
```

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

Specifies the volume number, 00 - FF.

```
X (for DS6000 and DS8000 models)
```

Specifies the address group, 0 - 1.

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.



```
SourceVolumeID:TargetVolumeID
```

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 DS8000 value for the XX is 00 - FE and the DS6000 value is 00 - 1F.

-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 (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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.

► setremoteflashrevertible -dev storage_image_ID -conduit LSS_ID -tgtoffline	
-tgtse	→
<pre>source_volume_ID:target_volume_ID</pre>	-▶◀

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 for the DS8000. The DS6000 value is 00 - 1F.

-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 (for a DS8000 model)
Specifies the logical subsystem number, 00 - FE.
```

- XY (for a DS6000 model) Specifies the logical subsystem number, 00 - 1E.
- ZZ (for DS6000 and DS8000 models) Specifies the volume number, 00 - FF.
- X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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.

1spprcpath

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.


Isavailpprcport

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 and DS6000 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 adaptors, and a newer release of the DS8000 (Release 6.0 or later), which supports 8-port host adaptors, the paths should connect only to the lower four ports on the newer storage unit.



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 unit, this command fails.

- S

(Optional). Displays the local port ID. You cannot use the -1 and the -s parameters together.

-1

(Optional). Displays all fields. You cannot use the -1 and the -s parameters together.

-remotedev storage_image_ID

(Required or Optional). Specifies the remote storage unit 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.

For DS8000, 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 for a DS8000 model and 00 - 1F for a DS6000 model.

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	Туре	Switch ID	Switch port
I0100	10200	FCP	N/A	N/A
I0150	10620	ESCON	N/A	N/A
10200	N/A	ESCON Switch	IBM.111.2222. 75113AB	I10
10250	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 unit. 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 unit. 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).

Ispprcpath

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 **-1** and the **-s** parameters together.

-1

(Optional) Displays the default output and the *Failed Reason* and *PPRC CG* descriptions. You cannot use the **-1** 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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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 **-1** parameter is used with the **1spprcpath** 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	

Src	Tgt	State	SS
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
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 unit.

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 unit.

The port ID component is 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 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 unit 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 units.

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 **-1** 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).

└──-tgtss── ss ID─┘ └─ -consistgrp─┘ └─ -forc	rce	

Parameters

Notes:

- 1. The **mkesconpprcpath** command is applicable **only** for the IBM Enterprise Storage Server (2105, Model 800 and Model 750). DS8000 models and DS6000 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. 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-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.

This is the only supported connectivity for machine types 2107 and 1750. Paths can be established between the following machine types: 2105:2105, 2107:2107, 2107:1750, 2107:2105, 1750:1750, 1750:2105.

Note: Secondary ports listed for PPRC, after issuing 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** issued 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, 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 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.

For DS8000, 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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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 only be used 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.

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 (z Systems)

-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 for the DS8000. The DS6000 value is in the range 00 - IF.

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.

For DS8000, 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]: YRemote 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.
- **1spprc** 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 unit 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.

▶failbackpprc	•
▶ failbackpprcdev storage_image_IDremotedev storage_image_ID	-
► -typemmir	•
►	•
► SourceVolumeID:TargetVolumeID	◄

Parameters

Notes:

- 1. You can issue 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 vice versa.
- **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 do not 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), 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.

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 which 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 unit that contains the target volume IDs that are defined by the **SourceVolumeID:TargetVolumeID** parameter. The **-remotedev** 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 **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 performed at a later time. The original order of updates is not 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 is not valid for a Global Copy (Extended Distance) remote mirror and copy volume relationship.
- 2. This parameter is not 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.

Note: This parameter applies only to z Systems volumes. It does not 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 is not 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 FlashCopy 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 are 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 allow you to 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 IBM Technical Support.

-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 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 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 multiple remote mirror and copy pair IDs with spaces.

A remote mirror and copy volume 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:

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0 - F for DS8000.

- XY (for a DS8000 model) Specifies the logical subsystem number, 00 - FE.
- *XY* (for a DS6000 model) Specifies the logical subsystem number, 00 - 1E.
- ZZ (for DS6000 and DS8000 models) 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.

For DS8000, 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.

►► failoverpprc	>
-dev— storage_image_IDremotedev— storage_image_ID_	-
► -type mmir	
► SourceVolumeID:TargetVolumeID	> 4

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 z Systems 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 FlashCopy pair between two volumes that have no previous relationship and only changed tracks are copied.

Notes:

- 1. This parameter can be used only as part of a Metro/Global Mirror (three-site) disaster recovery process.
- 2. 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 IBM Technical Support.

-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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0 - F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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.

For DS8000, 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.

Ispprc

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 **1spprc** 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 **-1** parameters together.

-1

(Optional) Displays the default output plus extra attributes that are identified as long output. You cannot use the **-s** and the **-1** 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(for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY(for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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 **1spprc** command

dscli> lspprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -1 0100:0100 0101:0101

The resulting output

ID	State	Reason	Туре	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 Sus- pended	Source LSS	Timeout (secs)	Critical Mode	First Pass Status	Incre- mental 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 **1spprc** 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. 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 Indicates 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 **1spprc** 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 **1spprc** command to view the out-of-sync value for a volume pair or a single volume on a DS8000 or a DS6000, 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. No value is displayed for a DS6000 model.

Date Suspended

Indicates the date when this relationship was last suspended. The value can be displayed as a " - ". No value is displayed for a DS6000 model.

SourceLSS

Indicates the consistency group LSS ID that is associated with the source volume of this PPRC volume relationship. No value is displayed for a DS6000 model.

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. No value is displayed for a DS6000 model.

Critical Mode

Indicates whether the remote copy and mirror primary volume represents a critical volume. No value is displayed for a DS6000 model.

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. No value is displayed for a DS6000 model.

Tgt Write

Indicates whether input is allowed to the remote mirror and copy secondary volume. No value is displayed for a DS6000 model.

GMIR CG

The value for this field is fixed to display N/A (not available). Use the **1ssession** 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 -1 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. No value is displayed for a DS6000 model.

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.

For DS8000, 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.

For DS6000, example: IBM.1750-75FA120

-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 z Systems^M 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 z Systems. The default value for this parameter is disabled.

Note: This parameter applies only to z Systems 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. Table 17 describes how the critical volume mode works.

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.
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.

Table 17. Critical volume mode definition overview

Notes:

- 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 unit; 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(for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY(for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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.

For DS8000, 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.1750-13AB79A -remotedev IBM.1750-13AB76A -type mmir -mode full -tgtread 1000:1205 1001:1206 1002:1207 1003:1208 1004:1209 and so on

 mkpprc -dev IBM.1750-13AB79A -remotedev IBM.1750-13AB76A -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.1750-13AB79A -remotedev IBM.1750-13AB76A -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 pmirsource_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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1 for DS6000 and 0 - F for DS8000.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

ZZ (for DS6000 and DS8000 models)

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.

		>
	- freezepprc	-
•	- Source_LSS_ID:Target_LSS_ID	

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.

For DS8000, 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.

For DS6000, example: IBM.1750-68FA150

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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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 (for a DS8000 model)
Specifies the logical subsystem number, 00 - FE.
```

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models) Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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.

For DS8000, 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 z System 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 S/390 or 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, 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 z Systems.

-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. The **-disableautoresync** parameter is available only in Version 5 Release 3 or later.

-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 (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

- *XY* (for a DS6000 model) Specifies the logical subsystem number, 00 - 1E.
- ZZ (for DS6000 and DS8000 models) Specifies the volume number, 00 - FF.
- X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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.

```
►►— rmpprc-
```

-dev— storage_image_ID_ _ -remotedev— storage_image_ID_ _ -at-

 $\begin{bmatrix} src \\ tgt \end{bmatrix}$



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.

For DS8000, 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.

For DS6000, example: IBM.1750-68FA150

-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 (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

- ZZ (for DS6000 and DS8000 models) Specifies the volume number, 00 - FF.
- X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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 unit 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.

▶▶—	unfreezepprc	
▶	Source_LSS_ID:Target_LSS_ID	►

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 unit 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.

For DS8000, 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.

For DS6000, example: IBM.1750-68FA120

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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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.

1sgmir 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, 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.

-lss LSS_ID

(Required) Specifies the master logical subsystem (LSS) that receives the **mkgmir** command. 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.

-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 failing the current consistency group. The drain time setting is required for a start action. If the drain time is not specified, the default drain time is 30 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 be started or resumed. 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 which a PPRC path has been established. An LSS ID is a two character value hexadecimal value in the following ranges:

- 00 FE, for a DS8000 model
- 00 1F, for a DS6000 model

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.

▶ pausegmirdev	storage_image_IDwithsecondarylssID	
► -session— session_ID-	<pre> — Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID</pre>	

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.

For DS8000, 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.

For DS6000, example of a fully qualified LSS ID: IBM.1750-68FA120/10

-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 for the DS8000. The DS6000 value is in the range 00 - 1F. 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.

DS8000 example of one Global Mirror association with a single subordinate in the configuration: IBM.2107-75FA120/00:IBM.2107-75FA150/00

DS8000 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

DS6000 example of one Global Mirror association with a single subordinate in the configuration: IBM.1750-68FA120/00:IBM.1750-68FA120/00

DS6000 example of multiple Global Mirror associations with two subordinates in the configuration: IBM.1750-68FA120/00:IBM.1750-68FA150/01 IBM.1750-68FA120/00:IBM.1750-68FA150/01

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.

For DS8000, 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 for the DS8000. The DS6000 value is in the range 00 - 1F.

For DS8000, 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. The following example is the same for a DS6000 with the exception of the storage image ID being different:

```
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 for the DS8000. The DS6000 value is in the range 00 - 1F. 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.

DS8000 example of one Global Mirror association with a single subordinate in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00

DS8000 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

DS6000 example of one Global Mirror association with a single subordinate in the configuration: IBM.1750-68FA120/00: IBM.1750-68FA150/00

DS6000 example of multiple Global Mirror associations with two subordinates in the configuration: IBM.1750-68FA120/00: IBM.1750-68FA150/00 IBM.1750-68FA120/00: IBM.1750-68FA150/01

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.

For DS8000, 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 **-1ss** 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.

For DS8000, 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 for the DS8000. The DS6000 value is in the range 00 - IF. 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.

DS8000 example of one Global Mirror association with a single subordinate in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00

DS8000 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

DS6000 example of one Global Mirror path association with a single subordinate in the configuration: IBM.1750-68FA120/00: IBM.1750-68FA150/00

DS6000 example of multiple Global Mirror path associations with two subordinates in the configuration: IBM.1750-68FA120/00: IBM.1750-68FA150/00 IBM.1750-68FA120/00: IBM.1750-68FA150/01

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 unit or on any of the subordinate storage units. The report that is generated by this command varies significantly depending on which storage unit that you issue the command and the parameters that you specify.

►►— showgmir

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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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 unit 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 unit 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 unit 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 unit 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 unit ID). All the other fields display as a "-" value.
- When you issue the **showgmir** command from the subordinate storage unit 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 unit.

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 Inter- val Time (sec- onds)	Coord. Time (milli- sec- onds)	Max CG Drain Time (sec- onds)	Current Time
IBM. 2107- 1300861 /14	1	0x25	Running	Not Fatal	0	50	30	10/25 /2006 13:45:44 PDT

CG Time	Succes- sful CG Percen- tage	Flash- Copy Sequ- ence Number	Master ID	Subor- dinate Count	Master/ Subordi- nate Assoc
10/25 /2006 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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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 unit.

Invoking the showgmir command

dscli> showgmir -dev IBM.2107-75FA120 -metrics 10

The resulting output

ID	Total Failed CG Count	Total Succes- sful CG Count	Succes- sful CG Percen- tage	Failed CG after Last Success	Last Succes- sful CG Form Time	Coord. Time (milli- seconds)	CG Inter- val Time (sec- onds)
IBM. 2107- 75FA120 /10	2	3	40	2	01/13 /1970 13:08:37 PST	50	5

Max CG Drain Time (sec- onds)	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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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.

Isgmir

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 **-1** parameters together.

-1

(Optional) Specifies that you want the system to display all the information that is associated with the session. You cannot use the **-1** 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 | -

(Optional) 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 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 **lsgmir** command using the **-l** parameter.

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	07/31/2009
					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.

For DS8000, 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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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.

For DS6000, example of a fully qualified LSS ID: IBM.1750-68FA120/10

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.

Showgmiroos ______ -dev—storage_image_ID____ -scope _____ si ____ -lss—LSS_ID—session_ID _____ *

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 for the DS8000 model. The DS6000 model value is in the range 00 - IF.

For DS6000, example of a fully qualified LSS ID: IBM.1750-68FA120/10

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 the **-dev** parameter is selected.

For DS6000, example: IBM.1750-75FA120

-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 for the DS8000. The DS6000 value is in the range 00 - IF.

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.

For DS6000, example of a fully qualified LSS ID: IBM.1750-68FA120/10

-action add | remove

(Required).

add Specifies that volumes be added to the session.

remove

Specifies that volumes be removed from the session.

-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 **-volume** parameter is required unless it is optionally replaced by the **-volpair** parameter.

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.

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:

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

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 share a common LSS ID.

The **-volpair** parameter is optional only as an alternative to the **-volume** parameter.

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.

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 0xXYZZ, 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

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
-lss 10 -action add -volume 1000-1010 01
```

The resulting output

Global Mirror session 01 successfully modified.

mksession

The **mksession** command opens a Global Mirror session.

▶ mksession	devs	storage_imag	ge_ID	remotedev	— storage_im	nage_ID	-lss— <i>ID</i> ——	>
► -volume—	volumo ID				lume_ID:targ		ID]
► Session	ID							> ∢

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.

For DS8000, 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.

For DS6000, example: IBM.1750-75FA120

-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 of 00 - FE for the DS8000. The DS6000 value is in the range 00 - 1F.

For DS6000, example of a fully qualified LSS ID: IBM.1750-68FA120/10

-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 (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models)

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.

Issession

The **Issession** 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.

For DS8000, example: IBM.2107-75FA120

- S

(Optional) Displays the session IDs. You cannot use the -1 and the -s parameters together.

-1

(Optional) Displays the default output. You cannot use the -1 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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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.

For DS6000, example of a fully qualified LSS ID: IBM.1750-68FA120/10

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 -1 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
IBM.2107- 75FA120 /10	02	Normal	IBM.2107- 75FA120 /1011	Active
IBM.2107- 75FA120 /10	02	Normal	IBM.2107- 75FA120 /1012	Remove Pending
IBM.2107- 75FA120 /10	02	Normal	IBM.2107- 75FA120 /1013	Join Pending

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
Primary Full Duplex	Secondary Simplex	True	Disabled
Prmary Simplex	Secondary Simplex	True	Disabled
Prmary Simplex	Secondary Simplex	False	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

LSSID	Session	Volume
10	01	1002
10	01	1003
10	02	1011
10	02	1012
10	02	1013

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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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 (for DS8000 models)

Specifies the address group, 0 - F.

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1F.


ZZ (for DS8000 models)

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.

For DS8000, 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 for the DS8000. The DS6000 value is in the range 00 - 1F.

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.

For DS6000, example of a fully qualified LSS ID: IBM.1750-68FA120/10

-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.

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.



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 = N996304B.tucson.ibm.com/9.11.178.201"
U,2005/10/04 15:29:37:432 MST,admin,1,,W,1001,User_Login_Expire,,0,
"IP = N996304B.tucson.ibm.com/9.11.178.201"
U,2005/10/04 15:32:56:979 MST,admin,1,,N,1000,User Login,,0,
"IP = N996304B.tucson.ibm.com/9.11.178.201"
U,2005/10/04 15:34:21:020 MST,admin,1,,N,1000,User_Login,,0,
"IP = N996304B.tucson.ibm.com/9.11.178.201"
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.
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.



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 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 devid value for the current command.

```
-auditlog | -sdocert | -etdata | -ipsec conn_ID[,conn_ID...] | -ipsecall | -ipsecstat |
-ipsecraw | -config
```

(Optional) You can choose only one of the following parameter options when you issue the command:

-auditlog

Exports one file containing the audit log for the management console server.

-sdocert

Exports the SDO certificate for each storage facility.

-etdata

Exports two files containing the IBM Easy Tier summary data.

-ipsec conn_ID[,conn_ID...]

Exports an IPSec connection file for each specified connection containing that connection's parameters.

-ipsecall

Exports all IPSec connections files containing connection parameters for each connection.

-ipsecstat

Exports the output of the strongSwan's **ipsec statusall** command.

Note: This parameter is currently available to assist with debugging IPSec connections; its availability and format are subject to change.

-ipsecraw

Exports the output of the last issued strongSwan **ipsec** command.

Note: This parameter is currently available to assist with debugging IPSec connections; its availability and format are subject to change.

-config

Offloads two files, one file containing the advanced settings, and another file containing the Install Corrective Service settings.

Note: Only one file set parameter must be specified.

-hmc 1 | 2 | all

(Optional) Specifies the HMC on which you want to export files. "-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.

-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 you to export files (using the **offloadfile** command) to a directory that requires elevated privileges. Unfortunately, the Windows operating system returns success and the **offloadfile** command displays a message stating that the files were exported successfully, but the files will not exist in the specified directory. To work around this problem, you can:

- 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.

Examples

Example 1

Invoking the **offloadfile** command to export EasyTier 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 has been offloaded to c:\temp\SF1300860ESS01_heat.data. Offloadfile: The etdata file has been offloaded to c:\temp\SF1300860ESS11 heat.data.

Example 2

Invoking the **offloadfile** command to export an SDO certificate.

dscli> offloadfile -dev IBM.2107-75FA120 -sdocert C:\temp

The resulting output

Date/Time: February 17, 2011 3:32:24 PM CST IBM DSCLI Version: 0.0.0.0 DS:IBM.2107-75FA120 CMUC00428I offloadfile: File for sdocert has been successfully copied to C:\temp\SD0Certificate_2107-941-1300860_2011-01-26_09_37_54.

Example 3

Invoking the **offloadfile** command to export audit logs

dscli> offloadfile
 -dev IBM.2107-75FA120 -auditlog C:\temp

The resulting output

Date/Time: February 17, 2011 3:32:24 PM CST IBM DSCLI Version: 0.0.0.0 DS: IBM.2107-75FA120 CMUC00428I offloadfile: File for auditlog has been successfully copied to c:\temp\IBM-2107-75FA120_auditlog_HMC1.log.

Example 4

Invoking the **offloadfile** command to export all IPSec connections dscli offloadfile -dev *IBM.2107-75FA120* -ipsecall C:\temp

The resulting output

Date/Time: February 17, 2011 3:32:24 PM CST IBM DSCLI Version: 0.0.0.0 DS: IBM.2107-75FA120

CMUC00428I offloadfile: File for ipsec has been successfully copied to C:\temp\IBM-2107-75FA120_conn_HMC1_MyConnection_1.conn. CMUC00428I offloadfile: File for ipsec has been successfully copied to

C:\temp\IBM-2107-75FA120_conn_HMC2_MyConnection_2.conn.

Chapter 5. Configuring access and security

DS CLI commands are used to configure access and security, such as Internet Protocol Security (IPSec).

Internet Protocol Security (IPSec) tasks

DS CLI commands are used to perform IPSec tasks.

Creating an IPSec connection configuration file

An IPSec connection configuration file is an ASCII text file that contains a connection definition. Use this procedure to create an IPSec connection configuration file.

You must have access to a text editor that can create ASCII text files to create a connection configuration file. You can create this file on the system where the DS CLI is installed, or on another system if the resulting connection file can be copied to the system where the DS CLI is installed.

An IPSec connection configuration file is an ASCII text file that contains a connection definition. Connection file formats are described at the beginning of this topic.

- You can use blank lines before and after the connection definition but not within the definition itself.
- The connection definition consists of a connection header followed by one or more connection attributes.
- The connection header consists of the keyword **conn** followed by the connection name. This name is used for all IPSec connection commands (except for the **mkipsec** command).
- The connection name is also used internally as part of a modified connection file name; therefore the connection name can be composed only of the following characters.
 - Upper and lower case alphabetic
 - Numeric
 - Special characters dash (-), underscore (_), and period (.)
- Each connection attribute consists of a keyword, an equal sign, and a keyword value. Space characters around the equal sign are optional.

The DS8000 system might support more attributes that are defined by strongSwan at http://wiki.strongswan.org/projects/strongswan/wiki/ConnSection. However, some of the strongSwan attributes are not allowed. For example, the attribute *keyexchange = ikev1* is not currently supported by the DS8000 system. Other attributes that were not specifically required for USGv6 testing might not have been tested.

Complete the following steps to create an IPSec connection configuration file.

- 1. Create a new ASCII text file using the extension .conn. This extension is not required, but it is a naming convention that is followed when you export connection files.
- 2. Enter the connection definition keyword **conn** to identify a connection definition block, followed by a blank space.
- **3**. Enter the connection name.
- 4. On subsequent lines, enter one or more attribute lines (with no blank lines). See all of the supported attributes in the IPSec command description.

Tip:

You can use the following attributes:

- Use %defaultroute for the **left** keyword. This keyword allows you to use the same connection file for multiple DS8000 servers that connect to a single end point, such as an encryption key server.
- Use %any for the **right** keyword. This keyword allows multiple remote end points to connect to the same DS8000 system without specifying a connection for each remote IP address.
- Use IP addresses instead of fully qualified domain names. This allows connections to be created even when DNS servers are not available.
- Specify explicit attributes (even when the default value is what the connection needs), since default values for specific keywords can change over time. Explicitly specifying the default attributes ensures that such changes do not affect your connections.
- 5. Save and close the file. Since the actual connection files that are used to create an IPSec connection are not saved on the HMC, it is important that you save a copy of the file. You can export the connection information from the DS8000 system into a file, but this does not ensure that the information that is exported is an exact match to the information that was in the original connection file.

The following lines are an example of a PSK IPSec connection configuration file.

```
conn connection1  # connection name is "connection1"
authby=psk  # deprecated keyword, but still accepted
auto=start
left=%defaultroute # equivalent to the HMC's IP address
right=9.12.212.17
type=tunnel
keyexchange=ikev2
esp=aes128-sha256
```

The following lines are an example of a public key IPSec connection configuration file.

```
conn connection2 # connection name is "connection2"
auto=start
left=%defaultroute # equivalent to the HMC's IP address
leftauth=pubkey
leftcert=certificateA.pem
right=9.12.212.17
rightauth=pubkey
rightcert=certificateB.pem
type=tunnel
keyexchange=ikev2
esp=aes128-sha256
```

Creating a PSK IPSec connection

Create a PSK (pre-shared key) IPSec connection by importing a connection file to the DS8000 system.

To create an IPSec connection, ensure that the following conditions are satisfied:

- You have a command-line interface prompt.
- You have a connection to a storage image that is to be used as one end point of an IPSec connection.
- The file that defines the IPSec connection is on the same server as the CLI.

The IPSec connection configuration file contains all of the specific information about the IPSec connection.

- 1. The **mkipsec** command imports this connection file to the DS8000 system and uses the information to create the IPSec connection in the IPSec server. The new connection might not be active depending on the connection file values, the IPSec server state, and the state of its other end point.
- After the new connection is created, the mkipsec command starts the IPSec server (if it is not running) or updates the IPSec server with the new connection information (if the IPSec server is already running). The IPSec server then has information about the connection, but the connection might not be active.

- If the **auto** attribute has the value "start", the IPSec server attempts to enable the connection. If the connection remote end point is enabled, the connection state becomes active.
- If the **auto** attribute has the value "add", the connection remains in the disabled state and you must manually enable it using the **chipsec** command.

The PSK connection protocol requires that both end points use the same pre-shared key. This key can be composed of any character string, including a hexadecimal string. You can specify the key directly or indirectly.

- Specify the key directly on the command line with the -phrase parameter.
- Type the key on a single line in an ASCII text file. (Limit access to the file through operating system access permissions.) You can then specify the file on the command line with the **-phrasefile** parameter.

If the **mkipsec** command fails, it returns an error message, attempts to restore the IPSec server to its state before the command was issued, and deletes the new connection information from the DS8000 system. The error message indicates the cause of the problem. After the problem is corrected, you can reissue the **mkipsec** command. If the error message is not clear, you can examine the strongSwan **ipsec** command raw output by using the **offloadfile –ipsecraw** command. The strongSwan **ipsec** command is the primary command that is used to control the IPSec server.

To create a PSK IPSec connection, complete the following steps.

1. Issue the **mkipsec** command.

Example 1: If you have the connection file C:\temp\connection1.conn and the pass phrase *my pass phrase*, enter the following command.

dscli> mkipsec -conn C:\temp\connection1.conn -phrase my pass phrase.

Example 2: To use the same connection file and the pass phrase file C:\Users\John\secrets\ psksecret.txt, enter the following command.

dscli> mkipsec -conn C:\temp\connection1.conn -phrasefile C:\Users\John\secrets\
psksecret.txt.

2. Issue the **lsipsec** command to view the status of the new connection.

Creating a public key IPSec connection

Use this procedure to create a public key IPSec connection by importing a connection file to the DS8000 system.

To create an IPSec connection, ensure that the following conditions are satisfied:

- You have a command-line interface prompt.
- You have a connection to a storage image that is used as one end point of an IPSec connection.
- You have a connection file that is on the same server as the command-line interface.

The process of creating a public key IPSec connection is similar to creating a PSK connection. Instead of using a pre-shared pass phrase, this process uses a public encryption key certificate for authentication.

- The DS8000 system cannot generate these certificates, so you must import existing certificates to the DS8000 system with the **mkipseccert** command.
- A connection can use the same certificate for both inbound and outbound communications, or it can use two different certificates.
- After the certificates are imported to the DS8000 system, the connection file is modified to add references to the certificates.
- The **-phrase** and **-phrasefile** parameters are not used because any phrases that are associated with the certificates are specified at the time the certificates are imported to the DS8000 system.

See an example of a public key connection file in the information about creating an IPSec connection configuration file.

To create a public key IPSec connection, complete the following steps.

Issue the mkipsec command to create a public key connection. For example, to use the C:\temp\connection2.conn connection file, enter the command as follows.

dscli> mkipsec -conn C:\temp\connection2.conn

2. Issue the **lsipsec** command to view the status of the new connection.

Enabling or disabling an IPSec connection

There might be circumstances where you would like to disable or enable a specific IPSec connection. The procedure for either process is identical except for the parameter that is used.

To manage an IPSec connection, ensure that the following conditions are satisfied:

- You have a command-line interface prompt.
- You have a connection to a storage image that is used as one end point of an IPSec connection.

Complete the following steps to disable or enable a DS8000 IPSec connection. In these steps, the connection *connection1* is an existing DS8000 IPSec connection.

1. Issue the **dscli** command to disable connection *connection1*. For example, enter the following command.

dscli> chipsec -disable connection1

- 2. Issue the **lsipsec** command to view the status of *connection1*.
- **3**. Issue the **dscli** command to enable the connection *connection1*. For example, enter the following command.

dscli> chipsec -enable connection1

4. Issue the **lsipsec** command to view the status of *connection1*.

Listing IPSec connections

Use this procedure to list information about one or all IPSec connections on a DS8000 system.

To list information about IPSec connections on your DS8000 system, ensure that the following conditions are satisfied:

- You have a command-line interface prompt.
- You have a connection to a storage image that is used as one end point of an IPSec connection.

If the DS8000 system cannot determine the status of a connection, you might be able to debug the problem by viewing the strongSwan **ipsec statusall** command output with the **offloadfile –ipsecstat** command.

To list information about IPSec connections on your DS8000 system, complete the following steps. In these steps, the connection *connection1* is an existing DS8000 IPSec connection.

1. Issue the **lsipsec** command to list information about any connection. For example, to list information about the connection *connection*1, issue the command as follows.

dscli> lsipsec connection1

Results similar to the following example can display.

2. Issue the **lsipsec** command to list information about all IPSec connections on your DS8000 system. For example, issue the command as follows.

dscli> lsipsec

Results similar to the following example display.

Export an IPSec connection file

You can use the **offloadfile** command to export a specific IPSec connection file or all connections files from a DS8000 system.

To export IPSec connection files from your DS8000 system, ensure that the following conditions are satisfied:

- You have a command-line interface prompt.
- You have a connection to a storage image that is used as one end point of an IPSec connection.

You can use the **offloadfile** command to export various files from a DS8000 system.

- Use the -ipsec parameter to export files from one or several IPSec connections. Separate the connection names with commas.
- Use the **-ipsecall** parameter to export all connection files.

Note: The original files that create a connection are not saved on the DS8000 system after the connection is created. Therefore, the exported IPSec connection files are not the same files that were used to create the IPSec connections.

- Specify a folder where the offloadfile command creates the exported connection files.
- Files that are created are named by the system with the storage image ID, the HMC ID, and the connection name.
- The **offloadfile** command does not overwrite any existing connection files by default. Specify the **-force** parameter to overwrite existing connection files.
- You can specify the **-conn** parameter with the **mkipsec** command to create a connection with the exported files.

To export IPSec connection files from your DS8000 system, complete the following steps. In these steps, the connection *connection1* is an existing DS8000 IPSec connection.

1. Issue the **offloadfile** command. to export the connection file. For example, issue the following command.

dscli> -ipsec connection1 C:\temp

2. Verify that a file, such as C:\temp\IBM.2107-75FA120_conn_HMC1_connection1.conn, was created.

Starting or stopping the IPSec server

Manual starting or stopping of the IPSec server is not required under normal conditions. Starting is automatic when the first connection is created. Stopping is automatic when the last connection is deleted. However, you can choose to start or stop the IPSec server manually under other conditions.

To start or stop the IPSec server, ensure that the following conditions are satisfied:

- You have a command-line interface prompt.
- You have a connection to a storage image that contains the IPSec server that you want to control.

The IPSec server does not run unless there is at least one IPSec connection that was created with the **mkipsec** command. Because the DS8000 system also uses IPSec connections for the remote VPN and call home functions, these functions also start and stop the IPSec server if no other connections were created with the **mkipsec** command.

- Stopping the IPSec server stops all existing active user connections. However, it fails if there are any active connections that are being used by the DS8000 system internally.
- Some connections might not reactivate when you restart the IPSec server, depending on the specifics of each connection. Therefore, do not manually stop the IPSec server unless directed to do so by IBM support personnel.

Note: Some of the connections that were created by the DS8000 system internally are not displayed with the **lsipsec** command.

You must use the **lsipsec** command to view the status of the server indirectly (you cannot directly display the state of the server).

- If the server stops, all connections display the state Server Down.
- If the server starts, each connection displays its state; no connection displays the state Server Down.

Complete the following steps to stop or start the IPSec server:

- Issue the setipsec command to stop the IPSec server. For example, issue the command as follows. dscli> setipsec –action disable –ctrl server
- 2. Issue the **lsipsec** command to indirectly view the server status.
- 3. Issue the **setipsec** command to start the IPSec server. For example, issue the command as follows. dscli> setipsec -action enable -ctrl server
- 4. Issue the **lsipsec** command to indirectly view the server status.

Importing certificates

Complete this task to import an encryption certificate (and a private key, if needed) to the DS8000 system.

To import a public key encryption certificate to your DS8000 system, ensure that the following conditions are satisfied:

- You have command-line interface prompt.
- You have connection to a storage image where you want to import the certificate.
- The connection file is on the same server as the CLI.

Public key cryptography usually uses a matched pair of asymmetrical keys (one public key and one private key) to send a message.

- 1. The sender encrypts the message with the public key and sends the encrypted message,
- 2. The recipient decrypts the message with the private key.

The public key is bound to the recipient's identity with a digital signature that is used to create an encryption certificate. The digital signature belongs to a certificate authority (CA), who verifies that the key and the identity in the certificate match. CA authentication ensures that the message can be decrypted only by a recipient who has the matching private key.

You can use the **mkipseccert** command to import both public and private encryption certificates. If a private key is required, it is also enclosed in a certificate that might require a pass phrase to access the private key. Since certificates are managed separately from the connections, you can use the same certificate for multiple connections. You can also update one certificate for all connections that use that certificate.

All certificates must use the PEM or DER format. The certificate ID that is used by the other IPSec certificate commands is the certificate file name that is specified by the **-cert** parameter (without the path information). The certificate name can contain only the following characters:

- Uppercase and lowercase letters
- Numbers
- Special characters dash (-), underscore (_), and period (.)

As a security measure, you can import the certificates, but you cannot export the certificates. Therefore, saving, protecting, and backing up the original certificate files are a user responsibility.

Complete the following steps to import an encryption certificate on the DS8000 system. This example assumes the following conditions:

- A public certificate C:\temp\ABC.pem.
- A private certificate C:\temp\ABCPrivate.pem.
- No pass phrase is required for the private key.
- 1. Issue the following command to import the certificates:

dscli> mkipseccert -cert "C:\temp\ABC.pem" -key C:\temp\ABCPrivate.pem

2. Issue the lsipseccert command to view the certificates on the DS8000 system.

Listing certificates

Use this procedure to list one or more encryption certificates on a DS8000 system.

To list a public key encryption certificate on your DS8000 system, ensure that the following conditions are satisfied:

- You have a command-line interface prompt.
- You have a connection to a storage image that contains the certificates that you want to list.

Because public and private certificates are a matched pair, only the name of the public certificate is displayed. The DS8000 system records the association between the two certificates and any pass phrase. If a private key certificate is associated with a public certificate, then is it is indicated in the HasKey column when you issue the **lsipseccert** command. If you delete a public certificate, the corresponding private certificate and pass phrase are also deleted.

Complete the following steps to list encryption certificates. The example assumes that there is an existing certificate named ABC.pem stored on the DS8000 system.

1. Issue the **dscli>lsipseccert ABC.pem** command to list the ABC.pem certificate. The displayed output can look similar to the following example.

2. Issue the **dscli>lsipseccert** command to list all certificates on the DS8000 system. The displayed output can look similar to the following example.

Chapter 6. Configuring and managing logical storage

You must complete a logical configuration for each storage system whether it consists of fixed block volumes or count-key-data volumes.

For example, each storage system must be assigned a worldwide node name (WWNN). You must also configure the arrays, ranks, logical subsystems or logical control units, and extent pools from which your logical volumes are created.

After your initial configuration, you might make adjustments in your configuration and, in time, you might delete your configuration and create a new one. The processes that are listed in this section are designed to help you complete these tasks by using the DS CLI commands.

Configuring new fixed block storage using the DS CLI

You can use the command-line interface to configure new fixed block storage within a storage image (or storage unit for DS6000).

Before you begin, you must be logged in to the DS CLI in interactive command mode. You must also be connected to a storage image (or storage unit for DS6000) that is used for open systems host system storage.

The creation of the fixed block storage configuration is described first. The configuration of the storage image (or storage unit for DS6000) SCSI host ports to enable access to fixed block storage is described second. You can run these two basic steps in the reverse order, but it is better to create storage configurations first, thus creating the media to back up configuration data that is not related to the storage configuration.

Configuring new fixed block storage involves the following processes:

- Creating fixed block extent pools
- Creating arrays
- · Creating and associating ranks with extent pools
- Creating fixed block volumes
- Creating fixed block volume groups
- Configuring Fibre Channel I/O ports
- Creating SCSI host port connections

Note: All the examples provided in the described tasks are based on the premise of using the interactive mode of DS CLI. If you were processing many transactions, you would likely use the script mode to process your transactions.

Creating extent pools for fixed block volumes using the DS CLI

Complete this task to create fixed block volume extent pools. This is the first step in configuring new fixed block storage. You can use the DS CLI commands to create extent pools for fixed block volumes.

Creating the extent pools before the arrays and ranks saves a processing step. When you create the ranks, you can assign them to existing extent pools. Otherwise, you must modify each rank object to complete the extent pool ID assignment after the extent pools have been defined.

Each extent pool is defined with the rank group of 0 or 1 and a storage type of **fb**. You must define one extent pool for each rank group and storage type combination. This means that you must make a minimum of two extent pools for a storage unit that contains fixed block storage: one fixed block extent pool per rank group.

Extent pools that are defined for rank group 0 or 1 are assigned an even- or odd-numbered extent pool ID, respectively. Even-numbered extent pools are managed by storage server ID 0. Odd-numbered extent pools are managed by storage server ID 1. Each rank is assigned to one extent pool; therefore, storage server workload is affected by the rank assignments to even- and odd-numbered extent pool IDs. It is better to evenly distribute rank and extent pool allocations to keep the storage server workloads balanced.

You can create more than the minimum number of extent pools. For example, you can define unique extent pools for each RAID type (5, 6 or 10) that is configured in a storage image. Or, you can define and name extent pools according to the host system attachments that access the volumes that are created from extent pool extents. You can have the same number of extent pools as ranks.

i5/OS considerations

i5/OS supports only specific volume sizes and these might not be an exact number of extents. i5/OS volumes are defined in decimal gigabytes (GB). You can use the following table when you are creating the logical volumes for use with i5/OS. You will notice that in almost every case, the i5/OS device size does not match a whole number of extents, so some space can be wasted for you specific configuration.

Pro- tected Model Type	Unpro- tected Model Type	i5/OS Device size (dec- imal giga- bytes)	Number of LBAs	Extents	Unus- able space (binary giga- bytes)	Usable space%
xxxx-A01	xxxx-A81	8.5	16 777 216	8	0.00	100.00
xxxx-A02	xxxx-A82	17.5	34 275 328	17	0.66	96.14
xxxx-A05	xxxx-A85	35.1	68 681 728	33	0.25	99.24
xxxx-A04	xxxx-A84	70.5	137 822 208	66	0.28	99.57
xxxx-A06	xxxx-A86	141.1	275 644 416	132	0.56	99.57
xxxx-A07	xxxx-A87	282.2	551 288 832	263	0.13	99.95

Note: Only Ax2, Ax4 and Ax5 models are supported as external LSU LUNs.

Use the **lsextpool** and **mkextpool** commands to create the fixed block extent pools. You must be logged into the DS CLI and connected to the storage unit that will be used for open systems host system storage.

Complete the following steps to create the fixed block extent pools. The example commands displayed in this task are shown in two formats. The first format shows the type of information the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Enter the **mkextpool** command to create the fixed block extent pool for rank group 0.

Example: dscli> mkextpool -dev *storage_image_ID* -rankgroup [0 | 1] -stgtype fb *extent_pool_name*

Example: dscli> mkextpool -dev IBM.2107-75FA120 -rankgrp 0 -stgtype fb P0 where P0 represents the extent pool name that you assign. This name can be 16 double-byte characters.

2. Press Enter. A successful process displays the following message: Extent pool P0 successfully created.

Note: The unique name that you assigned to the extent pool does not display in the process message. However, when you issue the **lsextpool** command, the extent pool name is displayed.

- **3**. Repeat Step 1 for each extent pool that you want to create. Try to evenly distribute rank and extent pool allocations to keep the storage server workloads balanced.
- 4. Verify the extent pool assignments by issuing the **lsextpool** command when you are done creating the extent pools. Use the **-1** parameter to display a full report for the extent pools that are assigned to the storage unit. Enter the **lsextpool** command at the dscli command prompt with the following parameters and variables:

Example: dscli> lsextpool -dev *storage_image_ID* -l

Example: dscli> lsextpool -dev IBM.2107-75FA120 -l

Creating arrays for fixed block volumes using the DS CLI

Complete this task to create arrays using the DS CLI commands.

The DS8000 machine type storage image storage devices (DDMs) are packaged into storage enclosure pairs. The DS8000 machine type contains at least one storage enclosure pair, with a minimum of 16 DDMs.

The DS6000 machine type contains at least one storage enclosure, with a minimum of four DDMs.

The DDMs of a storage enclosure are partitioned into array sites. The DS8000 machine type array site consists of eight DDMs, four from each storage enclosure of a storage enclosure pair, two-or-four (eight DDM) array sites per storage enclosure pair. The DS6000 machine type array site consists of four DDMs in one storage enclosure of a storage enclosure pair, with two-to-eight (four DDM) array sites per storage enclosure pairs must have identical capacity, rpm, and interface characteristics, and an interface to a common DA pair.

The creation of arrays is based on the array sites that are associated with the storage unit. Use the **lsarraysite** and **mkarray** commands to create the arrays.

You want to make an array from 1 or 2 array sites (DS6000) or 1 array site (DS8000). An array inherits the characteristics of its parent array sites and is given a RAID type attribute (5, 6, or 10). A DS8000 machine type array of RAID type 5, 6, or 10 is made from one (8 DDM) array site. A DS6000 machine type array object of RAID type 5 or 10 is made from one or two (4 DDMs) array sites. The status of the array is "unassigned" until the array is assigned to a rank.

Complete the following steps to create an array from unassigned array sites:

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Enter the **lsarraysite** command to view a list of array site IDs for all installed array sites. Review those arrays that are designated with the state of unassigned.

dscli> lsarraysite -dev storage_image_ID -state unassigned

Note: If this is your first time creating fixed block volumes, all the arrays are displayed with a state of unassigned.

- 2. Press Enter. A report of unassigned array sites is displayed. Use the list to identify unassigned array site capacity, rpm, and device adapter (DA) pair attributes. Record the RAID type for each array site.
- **3**. Enter the **mkarray** command to create an array from either one or two array sites (DS6000) or one array site (DS8000) with the status "unassigned".

dscli> mkarray -dev *storage_image_ID* -raidtype [5 | 6 | 10] -arsite *array_site*

Consider the following when you create the arrays:

For DS8000,

- Specify one array site with identical capacity, rpm, interface, and DA pair attributes.
- The new array inherits the capacity, rpm, interface, and DA pair characteristics of its parent array site.
- The state of the array remains unassigned until it is assigned to a rank.

For DS6000,

- Specify one or two array sites with identical capacity, rpm, interface, and DA pair attributes.
- The new array inherits the capacity, rpm, interface, and DA pair characteristics of its parent array sites.
- The state of the array remains unassigned until it is assigned to a rank.
- 4. Repeat Step 3 until all unassigned array sites have been assigned to an array.
- 5. Verify that the array-to-array site assignment is recognized and complete by issuing either the **lsarray** or **lsarraysite** command with the **-l** parameter.

Creating a rank using the DS CLI

Complete this task to create a rank using the DS CLI commands. A rank is a logically contiguous storage space that is made up of one array. You can assign a rank to every unassigned array.

A rank inherits the characteristics, including the RAID type, of its parent array and is given a storage type attribute of either FB (fixed block) or CKD (count key data). The rank configuration state is unassigned until it is assigned to an extent pool. An "unassigned" rank is not associated with either rank group 0 or 1. Any unassigned rank can be assigned to an extent pool that is associated with either rank group 0 or 1.

Note: You can assign a rank to an unassigned array and also assign the rank to an extent pool at the same time if you have already created the extent pools and the arrays. Creating extent pools first saves a step in the configuration.

Use the **lsarray**, **mkrank**, and **lsrank** commands to assign a rank to each unassigned array. You must be logged into the DS CLI and connected to the storage unit that will be used for open systems host system storage.

Note: You can enter the commands that are described in the steps either for a DS8000 model or for a DS6000 model. The storage image ID for the DS6000 model is different.

To create ranks, complete the following steps:

1. Ensure you have a list of the unassigned arrays for which ranks must be assigned. Enter the **1sarray** command to obtain this list if you do not already have it.

dscli> lsarray -dev IBM.2107-75FA120 -state unassigned

2. Enter the **mkrank** command to assign a rank to rank group 0 or 1 according to the rank group number of the assigned extent pool ID.

```
mkrank -dev IBM.2107-75FA120 -array A44 -stgtype fb
  -extpool P1
```

Notes:

- a. You can specify either the **-wait** or the **-extpool** parameter when you use the **mkrank** command. Either of these parameters allows you to be notified if the rank configuration has failed for any reason.
- b. If you use the **-wait** parameter, you cannot issue other commands until the entire transaction has processed.
- 3. Press Enter to display a report of rank assignments for your entire storage unit.

Because the process of creating the rank involves formatting drives, it could take some time before the process finishes. If you want to check on the process, enter the **1srank** command from a different DS CLI session. A successful process displays the following type of message:

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Device: IBM.2107-75FA120

Rank IBM.2107-75FA120/R44 successfully created.

- 4. Repeat Step 2 until all unassigned arrays are assigned a rank and an extent pool.
- 5. Enter the **lsrank** command to verify that ranks and extent pools have been assigned. dscli> lsrank -dev IBM.2107-75FA120 -1
- 6. Press Enter to display a report of the rank assignments for your entire storage unit.

Creating fixed block volumes using the DS CLI

Complete this task to create fixed block volumes.

Complete the following tasks before you create your fixed block volumes:

- Create your extent pools
- Create your arrays
- Create and assign your ranks

Complete the following steps to create fixed block volumes:

1. View your list of fixed block extent pool IDs and determine which extent pool IDs that you want to use as the source for the fixed block logical volumes. You obtained this list when you first created your extent pools. If this list is not available, issue the **lsextpool** command to obtain the list of extent pool IDs.

dscli> lsextpool -dev IBM.2107-13AAD7A -stgtype fb -l

Extent pool attributes determine the size and quantity of volumes that can be created. The extent pool ID (even/odd) indicates the storage server (0|1), which dictates that the LSS ID component of the volume ID must be an even or an odd number.

2. Enter the lsaddressgrp command to find unassigned and available address groups.

dscli> lsaddressgrp -dev IBM.2107-75FA120 -l

An address group refers to a group of LSSs. Up to 16 LSSs can be grouped into one address group. All LSSs in an address group must be of the same format (CKD or fixed block).

Note: If you are creating fixed block volumes for the first time, all the address groups are displayed with a state of "unassigned".

- **3**. Analyze the address group list to determine which LSSs can be used to make fixed block volumes. Consider the following conditions when doing your analysis:
 - If the address group list is empty, then all address groups are available to be defined (0 3).
 - If an undefined address group is used to create new fixed block volumes, select the lowest numbered address group.
 - If you are adding new fixed block volumes to an existing fixed block address group, use the **lslss** command to identify LSSs that are already defined in the target address group.
- 4. Enter the **mkfbvol** command to create fixed block volumes for the specified LSS.

dscli> mkfbvol -dev IBM.2107-75FA120 -extpool P1 -name finance#d -cap 8.6 0100-010F

Consider the following conditions regarding the command example in this step:

- All volumes can have the same type and capacity attributes.
- The **-extpool** parameter identifies a fixed block extent pool containing available data extents.
- The **-name** parameter allows you to assign an easy-to-use label or nickname to the volume. The volume name parameter can include a wildcard (#d or #h) that inserts a decimal or hexadecimal volume ID value into the volume name.

Note: The decimal designation does not apply to the volume ID number or the number of volumes that were created by the command. It only applies to the unique name that you have assigned. Also, when you process this command, the volume name that you have assigned does not appear in the confirmation message. To view the volume name that you have assigned, issue the **lsfbvol** or **showfbvol** command.

- The **-cap** (capacity) parameter is 8.6 GiB. The default is in gibibytes (GiB) where 1 GiB = 1 073 741 824 (2^30 bytes).
- The example provides a range of numbers (0100 010F) for the number of volumes to be created. Because volumes are created using the hexadecimal numbering system, the range in the example creates 16 volumes. The actual number of volumes that can be created is 255 per LSS based on the following criteria:
 - The volume ID is a 32 bit number that can be represented as four hexadecimal digits in the form
 of XYZZ where:

XY (for a DS8000 model)

Specifies the logical subsystem number, 00 - FE.

XY (for a DS6000 model)

Specifies the logical subsystem number, 00 - 1E.

ZZ (for DS6000 and DS8000 models) Specifies the volume number, 00 - FF.

X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

- For DS8000, you can define up to 255 LSSs in a storage unit. LSSs are either CKD or fixed block. Even numbered LSSs have an association with storage unit server 0. Odd numbered LSSs have an association with storage unit server 1.
- DS6000 has a 16 384 volume address space that is partitioned into 64 logical subsystem (LSS) units, where each LSS contains 256 logical volume numbers. The 64 LSS units are assigned to one of 4 address groups, where each address group contains 16 LSSs, or 4096 volume addresses. All of the LSSs in one address group must be of the same type (CKD or fixed block).
- 5. Repeat step 4 for all of the required logical volumes for each LSS.
- 6. Issue the **lsfbvol** command to display a report you can use to confirm the status of your fixed block volumes. Enter the **lsfbvol** command at the dscli command prompt with the following parameters and variables:

dscli> lsfbvol -dev IBM.2107-75FA120 -l -volgrp V2,V20

Note: The report can display that there was a configuration error that is associated with one or more of your **mkfbvol** transactions.

Creating LUN volumes for System i models

Complete this task to create fixed block LUN volumes for System i models.

Complete the following tasks before you create your fixed block LUN volumes:

- Created your extent pools
- 524 DS8000 Series Command-Line Interface User's Guide

- Created your arrays
- Created and assigned your ranks

When you begin your initial configuration, the LSSs and address groups do not exist. The LSSs are created when the first volume of the LSSs is defined during the processing of the **mkfbvol** command. The address group gets defined when the first LSS is defined.

When you create volumes, you must designate the logical subsystem (LSS) that a particular volume belongs to. After you assign a volume ID, use the first 2 digits to designate the LSS. For example, if you specify a volume ID of 1900, the volume then belongs to LSS 19.

Consider the following specifications before you create your fixed block LUN volumes:

- Volumes that belong to an even-numbered rank group (cluster) must be in an even-numbered LSS. Volumes that belong to an odd-numbered rank group (cluster) must be in an odd-numbered LSS. The cluster that a volume belongs to is determined by specifying the extent pool that the volume is assigned to.
- LSS numbers FF (DS8000) and 1F (DS6000) is reserved for internal use and must not be used as a volume ID.
- You must define each volume as protected or unprotected. This action is a notification to i5/OS; it does not mean that the volume is protected or unprotected. In reality, all LUNs are protected, either by RAID 5, RAID 6 or RAID 10. Defining a volume as unprotected means that it is available for i5/OS to mirror that volume to another internal or external volume of equal capacity. Unless you intend to use i5/OS (host-based) mirroring, define your logical volumes as protected.

Under some circumstances, you might want to mirror the i5/OS internal Load Source Unit (LSU) to a LUN in the DS6000 or the DS8000. In this case, define only one LUN volume as unprotected; otherwise, i5/OS attempts to mirror all unprotected volumes.

• In general, it is best to use one LSS for volumes from one rank.

Complete the following steps to create fixed block LUN volumes:

1. View your list of fixed block extent pool IDs and determine which extent pool IDs that you want to use as the source for the fixed block logical volumes. You obtained this list when you first created your extent pools. If this list is not available, enter the **lsextpool** command to obtain the list of extent pool IDs. Enter the **lsextpool** command at the dscli command prompt with the following parameters and variables:

dscli> lsextpool -dev IBM.2107-75FA120 -stgtype fb -l

Extent pool attributes determine the size and quantity of volumes that can be created. The extent pool ID (even \mid odd) indicates the storage server (0 \mid 1), which dictates that the LSS ID component of the volume ID must be an even or an odd number.

 Issue the mkfbvol command to create fixed block LUN volumes for the specified LSS. Enter the mkfbvol command at the dscli command prompt with the following parameters and variables: dscli> mkfbvol -dev IBM.2107-75FA120 -extpool p0 -os400 A05 -name i5_unprot_#h 1001-1002

Consider the following conditions regarding the command example in this step:

- The **-extpool** parameter identifies a fixed block extent pool that contains available data extents.
- The **-os400** parameter is available for designating the size and protection of a LUN volume by specifying the volume model. The example shows LUN volumes of protected model type A05 with a size of 35.1 decimal gigabytes (GB).
- The **-name** parameter is available for assigning an easy-to-use label or nickname to the volume. The volume name parameter can include a wildcard (#d or #h) that inserts a decimal or hexadecimal volume ID value into the volume name.

Note: The hexadecimal designation does not apply to the volume ID number or the number of volumes that were created by the command. It only applies only to the unique name that you

assigned. Also, when you process this command, the volume name that you assigned does not appear in the confirmation message. To view the volume name that you assigned, enter the **lsfbvol** or **showfbvol** command.

- The example provides a range of numbers (0101 0102) for the number of volumes to be created. The actual number of volumes that can be created is 255 per LSS based on the following criteria:
 - The volume ID is a 32-bit number that can be represented as four hexadecimal digits in the form of **XYZZ** where:
 - *XY* (for a DS8000 model) Specifies the logical subsystem number, 00 - FE.
 - XY (for a DS6000 model) Specifies the logical subsystem number, 00 - 1E.
 - ZZ (for DS6000 and DS8000 models) Specifies the volume number, 00 - FF.
 - X (for DS6000 and DS8000 models)

Specifies the address group, 0 - 1.

- You can define up to 255 (for DS8000) or 31 (for DS6000) LSSs in a storage unit. Even-numbered LSSs have an association with storage unit server 0. Odd-numbered LSSs have an association with storage unit server 1. LSS numbers FF (DS8000) and 1F (DS6000) are reserved.
- 3. Repeat step 2 for all of the logical volumes for each LSS.
- 4. Issue the **lsfbvol** command to display a report you can use to confirm the status of your LUN volumes. Enter the **lsfbvol** command at the dscli command prompt with the following parameters and variables:

dscli> lsfbvol -dev IBM.2107-75FA120 -l

Correcting a fixed block configuration error

Complete this task to correct a fixed block volume configuration error.

There might be occasions when you are using the **mkfbvol** command to create fixed block volumes, but the transaction fails. You might not be aware of the failure until you run the **lsfbvol** or the **showfbvol** command to check the status of the volumes that you created.

The **lsfbvol** or the **showfbvol** commands display reports that include a *configstate* category. The configuration state category reports on the current state of the volume. One of the configuration state codes is **configuration error**.

A status of **configuration error** specifies that the configuration process did not complete successfully. This state reflects an internal error condition and is not an indication that there was a user input error.

You might want to gather additional information about what caused the error, which can help you determine how to correct it. To correct this error state, you must delete the designated volume configuration and submit a new transaction request.

Complete the following steps to obtain additional information about the configuration error and to correct this error condition.

1. Add the **-v** (verbose) command parameters to your **mkfbvol** command, and reissue the command for the transactions that show the configuration error designation.

Note: You can also turn on the verbose mode in your profile file, and reissue the command. If you designate the verbose mode, the display of extra output includes the error code that is generated when the create volume transaction fails.

2. Issue the **rmfbvol** command to delete the designated volume configurations if you do not want to obtain additional information about what caused the configuration error.

Note: In most of instances, this method is the only one for correcting a configuration error.

Creating fixed block volume groups using the DS CLI

Complete this task to create fixed block volume groups.

A volume group identifies the set of fixed block logical volumes that are accessible by one or more SCSI host system ports. SCSI host system access is constrained to the identified access mode. Only those SCSI host ports that are registered to a volume group ID are allowed to access the set of logical volumes that is contained by the volume group.

Logical volumes can be assigned to a volume group when the volume group is created, or the logical volumes can be added (or removed) later. The volume group type determines the maximum number of volumes that can be assigned to a volume group, either a maximum of 256 volumes or a maximum of 64 000 volumes. The volume group type must be selected according to the addressing capability of the SCSI host system that uses the volume group.

Complete the following steps to create and view fixed block volume groups:

1. Issue the **mkvolgrp** command to create a fixed block volume group. Enter the **mkvolgrp** command at the dscli command prompt with the following parameters and variables:

Note: Repeat this step for each volume group that you want to create. dscli> mkvolgrp -dev IBM.2107-75FA120 -hosttype pSeries -volume 0001-0010,0120 my_nickname

Notes:

- a. You can use the -hosttype parameter with the mkvolgrp command. This parameter is an easier way of specifying the type of volume group. If you do not use the -hosttype parameter, it is assumed that the volume group type is *scsimask*.
- b. You cannot use the **-type** parameter and the **-hosttype** parameter together.
- c. If your volume group is not *scsimask* type and you do not want to use the **-hosttype** parameter, use the **-type** parameter. *scsimask* as the default value of the **-type** parameter; you can also specify *scsimap256* or *os400mask* as your volume group type. For information about the criteria that is associated with these volume group types, see the **mkvolgrp** command.
- d. Volume IDs must meet the following criteria:
 - ID ranges must be separated by a comma (displayed as 0001-0010,0120 in the example).
 - For *scsimap256*, the array or ranges cannot exceed 256 volume ID entries. Otherwise, up to 64 384 entries are allowed.
 - Use the **-type** *0s400mask* parameter if the volume group is limited to fixed block volume OS400-protected or OS400-unprotected types. Otherwise, the volume group is limited to the fixed block volume type DS8000 or DS6000 machine type.
 - The volume group name (*my_nickname* in the example command) must be unique within the scope of the specified storage image.
- Issue the lsvolgrp command to create a list of assigned volume group IDs. Enter the lsvolgrp command at the dscli command prompt with the following parameters and variables: dscli> lsvolgrp -dev IBM.2107-75FA120 -1

Notes:

- a. The **lsvolgrp** command with the *-l* parameter displays a report with the following three values:
 - Name (the unique name that you assigned to the volume group)

- Volume group ID
- Type (the configured volume group type)
- b. You can narrow the scope of the report by requesting a specific type of volume. See the **lsvolgrp** command for information about the **-type** parameter.

Creating a volume group for System i models

Complete this task to create volume groups for System i models so that volumes can be assigned to the i5 Fibre Channel adapters.

Use the **mkvolgrp** command to create a volume group that contains the volumes to be assigned to an i5 Fibre Channel adapter. The following considerations determine how you will create your volume groups:

- If you are using a multipath connection, a volume group is assigned to two or more i5 Fibre Channel adapters. Each Fibre Channel adapter provides one path to volumes in the volume group.
- If you are using an external load source, you create a volume group that contains one volume. After a partition is initially loaded from the external load source, you can add more volumes to this volume group so that the i5/OS recognizes them and can use them.

When you create a volume group for i5/OS, you should specify the **-hosttype iSeries** parameter as part of the **mkvolgrp** command. The **-hosttype iSeries** parameter saves some processing time because this parameter automatically supplies information that would have to be specified separately. For example, the i5/OS uses a logical blocksize of 520, and volumes that are created for i5/OS use a blocksize of 520 bytes. By specifying the **-hosttype iSeries** parameter, you also denote that the logical size of the blocks in the volumes is 520 bytes.

Complete the following steps to create and view volume groups for System i models:

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **mkvolgrp** command to create a volume group. Enter the **mkvolgrp** command at the dscli command prompt with the following parameters and variables:

Note: Repeat this step for each volume group that you want to create.

The following two examples provide the commands that you can use to create volume groups depending on whether you use the external load source. The first example creates a volume group that contains one unprotected volume if you do use an external load source. (If you are using an external load source, you can initially only have one volume in the volume group.) The second example creates a volume group that contains all volumes if you do not use an external load source.

Example 1 (using an external load source)

dscli> mkvolgrp -dev IBM.2107-13ABVDA -hosttype iSeries -volume 1000 blue CMUC00030I mkvolgrp: Volume group V14 successfully created.

Example 2 (not using an external load source)

dscli> mkvolgrp -dev IBM.2107-13ABVDA -hosttype iSeries -volume 1000-1002 blue CMUC00030I mkvolgrp: Volume group V14 successfully created.

Notes:

- a. The confirmation message for the end process shows that the created volume group is automatically assigned an ID that is different from the name of the volume group that you specify in the command. You will see the name that you assigned associated with the volume group when you use the **lsvolgrp** and **showvolgrp** commands. However, if you want to work specifically with the volume group, you must reference the volume group ID.
- b. This volume group is also referred to as SCSI520-MASK. When an error message is displayed for the OS400 MASK, SCSI520-MASK is referenced instead.

- **c.** If you do not use an external load source, it is a good practice to create a volume group that contains all the volumes that will be assigned to the i5 Fibre Channel adapter.
- d. Some System i models only support 32 device addresses per volume group.
- Issue the lsvolgrp command to create a list of assigned volume group IDs. Enter the lsvolgrp command at the dscli command prompt with the following parameters and variables: dscli> lsvolgrp -dev IBM.2107-13ABVDA -1

Notes:

- a. The **lsvolgrp** command with the **-l** parameter displays a report with the following three values:
 - Name (the unique name that you assigned to the volume group)
 - Volume group ID (the identification number of the volume group)
 - Type (the configured volume group type)
- b. You can narrow the scope of the report by requesting a specific type of volume. See the **lsvolgrp** command for information about the **-type** parameter.
- **3**. Verify your host type information by issuing the **lshosttype** command using the following command format at the dscli command prompt:

dscli> lshosttype -type os400mask

This command displays a report like the following example:

Name	Profile	AddrDiscovery	LBS
iSeries	IBM iSeries - os400	reportlun	520

Note: You can obtain the same results if you use the -type os400all parameter.

Configuring Fibre Channel I/O ports using the DS CLI

Complete this task to configure Fibre Channel I/O ports using the DS command-line interface.

Before you begin, you must have the command-line interface prompt, and you must be connected to a storage image that is to be used for open systems host system storage.

In this process, you must designate the topology for the I/O port. The following three topology settings are available:

fibre channel arbitrated loop (coded as fc-al in the setioport command)

Enables the SCSI ULP with a FC-AL topology. The FC-AL topology does not support PPRC path I/O operations.

SCSI-FCP (coded as scsi-fcp in the setioport command)

Enables the SCSI ULP with a point-to-point or switched fabric topology. PPRC path I/O operations are enabled for this setting.

Note: Designate this topology for System i models using i5/OS level V5R3M5 and above.

ficon (coded as ficon in the setioport command)

Enables the FICON ULP with a point-to-point or switched fabric topology. PPRC path I/O operations are not supported for FICON ULP.

The storage image supports the Fibre Channel host adapter (HA) card type. For DS8000 machine types, HA cards are installed in I/O enclosures, each containing up to four HA cards. For DS6000 machine types, one or two HA cards are installed in each of the two processor node assemblies. Use the **lsioport** and **setioport** commands to configure Fibre Channel I/O ports.

Each Fibre Channel HA card contains four I/O ports. The storage image microcode automatically creates one I/O port to represent each HA card I/O port. The default Fibre Channel I/O port settings enable SCSI-FCP "identified" access to fixed block volumes. You might have to modify the I/O port settings to enable SCSI FC-AL access to fixed block volumes.

To configure the Fibre Channel ports, complete the following steps:

1. View a list of Fibre Channel port IDs by typing the following command format at the dscli command prompt:

dscli> lsioport -dev IBM.2107-75FA120 -l -type fc

A detailed report is displayed that lists the Fibre Channel I/O ports.

2. Analyze the report and determine which I/O port IDs that you want to access the fixed block volumes.

Configure a minimum of four I/O ports for SCSI host I/O operations. Select ports with physical locations on different host bus adapter (HA) cards. If possible, locate the HA cards in different I/O enclosures.

3. Set the I/O ports that you have identified to enable the FC-AL (Fibre Channel arbitrated loop), SCSI-FCP, or FICON topology. The following example shows how to enable the FC-AL topology by typing the following command format at the dscli command prompt:

Note: I/O ports are automatically set to the offline state and returned to the online state after configuration changes are applied.

dscli> setioport -dev IBM.2107-75FA120 -topology fc-al 0012 0013 0112 0113

4. Press Enter. A successful process returns a confirmation message indicating that the port IDs have been successfully configured.

Creating SCSI host port connections using DS CLI

Complete this task to create SCSI host port connections using the DS command-line interface.

Before you begin, you must have the command-line interface prompt, and you must be connected to a storage image that can be used for open systems host system storage.

The DS8000 and DS6000 machine types support the "identified" access mode for SCSI host attachments, which requires that all SCSI host ports be identified to a DS8000 storage image or DS6000 storage unit respectively. One SCSI host port connection must be created for each SCSI host port that accesses DS8000 storage image and DS6000 storage unit fixed block volumes. Use the **lshosttype**, **mkhostconnect**, and **lshostconnect** commands to create the SCSI host port connections.

A SCSI host port contains attributes that identify the following information:

- SCSI host system type
- Port profile
- Port WWPN
- Volume group ID that the port accesses
- An array of storage unit or storage image I/O port IDs that the host port logs into for volume access
- An attribute to indicate that all I/O ports can be used for volume access
- Host port description
- Port nickname

There are two ways that you can approach this task:

• Use the **-hosttype** parameter with the **mkhostconnect** command. Using the **-hosttype** parameter is the best solution for most users.

• Use the **mkhostconnect** command with the **-lbs**, **-addrdiscovery**, and **-profile** parameters.

Notes:

- 1. Specifying the -hosttype parameter automatically sets the -lbs, -addrdiscovery, and -profile values.
- 2. If you do not use the **-hosttype** parameter, you must issue the **lsportprof** command to ensure that you obtain the correct values to use with the **-lbs**, **-addrdiscovery**, and **-profile** parameters.
- 3. You cannot use the **-hosttype** parameter with these other parameters.

The following task is described from the assumption that you have used the **-hosttype** parameter.

To configure the SCSI host ports, complete the following steps:

Obtain your host type information by issuing the lshosttype command. Enter the lshosttype command at the dscli command prompt with the following parameters and variables:

dscli> lshosttype -l -type volumeGroup_type

This command displays a report like the following example:

Name	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)
iSeries (if os400all was specified)	IBM iSeries - os400	reportlun	520	IBM iSeries Servers (System i)

Note: Volume group type is one of the following designations (use a separate command for each choice):

- ficonall
- scsiall
- scsimask
- scsimap256
- os400all
- os400mask

The same results are displayed when you specify os400all or os400mask or when you specify scsiall and scsimask or scsimap256.

2. Create SCSI host ports by issuing the **mkhostconnect** command. Enter the **mkhostconnect** command at the dscli command prompt with the following parameters and variables:

```
dscli> mkhostconnect -dev storage_image_ID -wwpn wwpn
    -hosttype host_type -volgrp volume_group_ID -ioport port_ID
    host_name
```

Notes:

- a. The **-wwpn** parameter specifies the 16-character worldwide name that is assigned to the host system Fibre Channel adapter port. This WWPN value is validated each time that the host system port logs into an I/O port.
- b. The **-hosttype** parameter specifies Fibre Channel communications layer characteristics that might be unique according to the host system manufacturer, operating system, or version of the system. Typical specifications are iSeries, pSeries, an so on.

- **c.** The **-volgrp** parameter specifies the volume group ID that this host port can access. Host port objects might be created prior to creating volume groups, in which case you must use the **chhostconnect** command to add volume group ID assignments at a later time.
- d. The **-ioport** *all* specifies SCSI host port (WWPN) access into all IO ports that are configured for the FC-AL or SCSI-FCP topology.
- e. host_name specifies the SCSI host system nickname that you have assigned.
- 3. Repeat Step 2 for each SCSI host system port that will access LUN volumes.
- 4. Verify that all SCSI host ports have been configured and that they are recognized by the storage unit according to your specifications by issuing the **lshostconnect** command with the **-l** parameter.

Configuring new count key data storage using DS CLI

Your storage image for a z Systems host system must be configured for new count key data (CKD) storage.

Before you begin, you must be logged into the DS CLI in interactive command mode. You must also be connected to a storage unit that is used for z Systems models host storage.

Configuring CKD storage involves two basic processes: the creation of the CKD storage configuration and the configuration of the storage unit I/O ports for z Systems host system attachment. These two basic processes can be performed in the reverse order, but it is better to create storage configurations first. Creating the storage configuration first creates the media to back up configuration data that is not related specifically to the storage configuration.

Configuring new CKD storage involves the following processes:

- · Creating CKD extent pools
- Creating arrays
- Creating and associating ranks with extent pools
- Creating logical control units
- Creating CKD volumes
- Creating CKD volume groups (system generated).

The internal microcode automatically creates the CKD FICON/ESCON All volume group ID (V10) for DS8000, and it creates the CKD FICON All volume group ID (V10) for DS6000, and automatically assigns all CKD base and alias volumes to this volume group. This volume group ID (V10) is automatically assigned to storage unit I/O ESCON ports, and to I/O Fibre Channel ports (for DS8000), and to I/O Fibre Channel ports (for DS6000) that are configured for FICON I/O operations. For DS8000, the ESCON I/O ports are constrained to access Address Group 0 volume IDs (0000-0FFF).

• Configuring Fibre Channel I/O ports

Creating count key data extent pools using the DS CLI

You can use DS CLI commands to create extent pools for CKD volumes. This step is the first in configuring new count key data storage.

You must be logged in to the DS CLI and connected to the storage unit that is to be used for your z Systems host system to complete this task.

You can create the extent pools before the arrays and ranks to save a processing step. When you create new ranks, you can assign them to existing extent pools. Otherwise, you must modify each rank object to complete the extent pool ID assignment after the extent pools have been defined.

Each extent pool is defined with the rank group of 0 or 1 and a storage type of *ckd*. You must define at least one extent pool for each rank group and storage type combination. This requirement means that you must make a minimum of two extent pools for a storage unit that contains CKD storage: one CKD extent pool per rank group.

Extent pools that are defined for rank group 0 or 1 are assigned an even- or odd-numbered extent pool ID, respectively. Even-numbered extent pools are managed by storage server ID 0. Odd-numbered extent pools are managed by storage server ID 1. Each rank group is assigned to one extent pool; therefore, storage server workload is affected by the rank assignments to even- and odd-numbered extent pool IDs. To keep the storage server workloads balanced, evenly distribute rank and extent pool allocations.

Notes:

- 1. You can create more than the minimum number of extent pools. For example, you can define unique extent pools for each RAID type (5, 6 or 10) that is configured in a storage image. Or, you can define and name extent pools according to the host system attachments that access the volumes that are created from extent pool extents.
- 2. You can have the same number of extent pools and ranks.

Complete the following steps to create the extent pools:

 Issue the lsextpool command to display a list of the existing CKD extent pools. The following line is an example of the command that you can type in the dscli command prompt: lsextpool -dev *IBM.2107-75FA120* -stgtype ckd

where IBM.2107-75FA120 is the name of the storage image ID.

- 2. Analyze the extent pool listing for the following information:
 - Does the minimum set of extent pools exist? There must be one extent pool for rank group 0 and one extent pool for rank group 1.

Note: If extent pools are being created for the first time, the minimum number of extent pools does not exist.

• Does each extent pool have a rank group that is assigned to it and are they balanced?

Note: If extent pools are being created for the first time, there are no rank assignments.

- Are additional extent pools required?
- **3**. Issue the **mkextpool** command to create the extent pools. The following examples show the commands that you can type in the dscli command prompt:

Remember: You must create a minimum of two extent pools. There must be one extent pool for rank group 0 and one extent pool for rank group 1.

```
mkextpool -dev IBM.2107-75FA120 -rankgrp 0 -stgtype ckd
extent_pool_name
mkextpool -dev IBM.2107-75FA120 -rankgrp 1 -stgtype ckd
extent_pool_name
```

where *extent_pool_name* is the name that you want to assign to the extent pool. This parameter is a required. The name cannot be longer than 16 characters.

Create additional extent pools for each of the following conditions:

- Each RAID type (5, 6 or 10)
- Each disk drive module (DDM) size
- Each CKD volume type (3380, 3390)
- Each logical control unit (LCU) address group
- 4. Press Enter.

Note: The unique name that you assigned to the extent pool does not display in the process message. However, when you issue the **lsextpool** command, the extent pool name is displayed. The following line is an example of the message that is displayed for a successful process: Extent pool P1 successfully created.

5. Repeat step 2 on page 533 for each extent pool that you want to create.

Remember: To keep the storage server workloads balanced, evenly distribute rank and extent pool allocations.

6. After you have created the extent pools, issue the **1sextpool** command to verify the extent pool assignments. Use the **-1** parameter to display a full report for the extent pools that are assigned to the storage unit. The following line is an example of the command that you can type in the dscli command prompt:

lsextpool -dev IBM.2107-75FA120 -1

7. Optionally, print the list of extent pools so that you can refer to this list when you create CKD volumes.

Creating arrays for CKD volumes using the DS CLI

Complete this task to create arrays for CKD volumes using the DS CLI commands.

The DS8000 machine type storage image storage devices (DDMs) are packaged into storage enclosure pairs. The DS8000 machine type contains at least one storage enclosure pair, with a minimum of 16 DDMs. All DDMs that are installed in a storage enclosure pair have identical capacity, rpm (revolutions per minute), and interface characteristics.

The DS6000 machine type must contain at least one storage enclosure, with a minimum of four DDMs.

The DDMs of a storage enclosure are partitioned into array sites. A DS8000 machine type array site consists of eight DDMs, four from each storage enclosure of a storage enclosure pair, two-or-four (eight DDM) array sites per storage enclosure pair. A DS6000 machine type array site consists of four DDMs in one storage enclosure of a storage enclosure pair, with two to eight (four DDM) array sites per storage enclosure pair. All array sites of a storage enclosure pair have identical capacity, rpm, and interface characteristics, and an interface to a common DA pair.

The creation of arrays is based on the array sites that are associated with the storage unit. You must make an array from 1 or 2 array sites. An array inherits the characteristics of its parent array sites, and is given a RAID type attribute (5, 6, or 10). A DS8000 machine type array object of RAID type 5, 6, or 10 is made from one (8 DDMs) array site. A DS6000 machine type array object of RAID type 5 or 10 is made from one or two (4 DDMs) array sites.

Note: The array status is "unassigned" until the array is assigned to a rank.

Use the **lsarraysite** and **mkarray** commands to create the arrays. You must be logged into the DS CLI and connected to the storage unit that will be used for open systems host system storage.

Complete the following steps to create arrays for a CKD volume configuration:

 Issue the lsarraysite command to find the unassigned array sites. Enter the lsarraysite command at the dscli command prompt with the following parameters and variables: dscli> lsarraysite -dev storage image ID -state unassigned

Note: If this is your first time creating volumes, you will see all the arrays with a state of "unassigned".

2. Press Enter. A report of unassigned array sites is displayed. Use the list to identify unassigned array site capacity, rpm, and device adapter (DA) pair attributes. Record the RAID type for each array site.

3. Issue the mkarray command to create an array from each site with the status "unassigned". Enter the mkarray command at the dscli command prompt with the following parameters and variables: dscli> mkarray -dev storage_image_ID -raidtype [5 | 6 | 10] -arsite array_site Repeat this command until all unassigned array sites have been assigned to an array.

Notes:

- **a.** You can specify one or two array sites for DS6000. If there are two array sites, both sites must be associated with a common DA pair ID. Two array sites must be separated by commas with no blank space in between. Example: S10,S11.
- b. The new array site inherits the capacity, rpm, interface, and DA pair characteristics of its parent array site or sites depending on the machine type. The state of the array is "unassigned" until it is assigned to a rank.

Creating a rank for CKD volumes using the DS CLI

Complete this task to create a rank for a CKD volume. A rank is a logically contiguous storage space that is made up of one or more arrays. You want to assign a rank to every unassigned array.

A rank inherits the characteristics, including RAID type, of its parent array and is given a storage type attribute FB (fixed block) or CKD (count key data). The rank configuration state is unassigned until it is assigned to an extent pool. An unassigned rank is not associated with either rank group 0 or 1. Any unassigned rank can be assigned to an extent pool that is associated with either rank group 0 or 1.

Note: You can assign a rank to an unassigned array and assign the rank to an extent pool at the same time if the extent pools and the arrays were created previously. Creating extent pools first saves a step in the configuration.

Use the **lsarray**, **mkrank**, and **lsrank** commands to assign a rank to each unassigned array. You must be logged in to the DS CLI and be connected to the storage unit that is used for open systems host system storage.

To make ranks, complete the following steps:

1. Enter the **lsarray** command to ensure that you have a list of the unassigned arrays for which ranks must be assigned.

dscli> lsarray -dev IBM.2107-75FA120 -state unassigned

2. Enter the **mkrank** command to assign a rank to rank group 0 or 1 according to the rank group number of the assigned extent pool ID.

```
dscli> mkrank -dev IBM.2107-75FA120 -array A44
-stgtype ckd -extpool P1
```

Notes:

- a. You can specify either the **-wait** or the **-extpool** parameter when you use the **mkrank** command. Either of these parameters generates output if the rank configuration fails for any reason.
- b. When you use the **-wait** parameter, you cannot enter any other commands until the entire transaction is done processing.
- **3**. Press Enter to create the ranks.

The process of making the rank involves formatting drives. It can take a little time before the process finishes. To check on the process, enter the **lsrank** command from a different DS CLI session. A successful process generates the following type of message:

Rank IBM.2107-75FA120/R44 successfully created.

- 4. Repeat Step 2 and step 3 until all unassigned arrays are assigned a rank and an extent pool.
- 5. Enter the **lsrank** command to verify that ranks and extent pools are assigned. dscli> lsrank -dev IBM.2107-75FA120 -1

6. Press Enter. A report of the rank assignments for your entire storage unit is displayed.

Creating logical control units for CKD volumes using DS CLI

The logical control unit (LCU) is the z Systems host equivalent of the logical subsystem (LSS) for open systems hosts. The LCU must be defined (created) before CKD logical volumes can be created.

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.

The DS6000 has a 16 384 volume address space that is partitioned into 64 logical subsystem (LSS) units, where each LSS contains 256 logical volume numbers. The 64 LSS units are assigned to one of 4 address groups, where each address group contains 16 LSSs, or 4096 volume addresses. All of the LSSs in one address group must be of the same type (CKD or fixed block).

Typically, LCUs are created in groups of 16, beginning at LSS address X'x0'.

Use the **lsaddressgrp**, **mklcu**, and **lslcu** commands to create the LCU type logical subsystems. You must be logged into the DS CLI and connected to the storage unit that will be used for open systems host system storage.

To create LCUs, complete the following steps:

 Find unassigned and available address groups by issuing the lsaddressgrp command. To use the lsaddressgrp command, type the following at the dscli command prompt: dscli> lsaddressgrp -dev IBM.2107-75FA120

This command displays a report on the status of the address groups within your storage unit.

- **2**. Analyze the report to identify all of the address groups that are available to be defined. Use the following criteria:
 - If the list is empty, all of the address groups are available to be defined.
 - A defined address group with the storage type fb (fixed block) is not available to be defined.
 - A defined address group with the storage type ckd and with fewer than 16 LSSs is available for LCU definition.
 - If you are using an undefined address group to make new LCUs, select the lowest numbered address group that is not defined.
 - If you are defining a new LCU in an existing CKD address group, use the **lslcu** command to identify LCUs that are already defined in the target address group.
- **3**. Make the LCU logical subsystem objects by issuing the **mklcu** command. Type the command using the following format at the dscli command prompt:

dscli> mklcu -dev IBM.2107-75FA120 -qty 16 -id 00 -ss 0010 -lcutype 3390-3

In this example, the values specify the following information:

qty

Specifies the number of LCU IDs to be created.

- id Specifies the LCU ID to be created, or the first LCU ID in a sequence of LCU IDs to be created.
- **ss** Specifies the subsystem ID that you have assigned. 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.

lcutype

Specifies the type of LCU to be created. You can specify the following types:

- 3390-3
- 3990-tpf

- 3990-6
- bs2000
- 4. Press Enter. A successful process displays a confirmation message listing each LCU ID number that has been successfully created.
- 5. Verify that the LCUs are recognized in the storage unit by issuing the **lslcu** command at the dscli command prompt as follows:

dscli> lslcu -dev IBM.2107-75FA120 -l

Using the **-1** parameter displays a more detailed report for each LCU that is associated with your storage unit.

Creating count key data volumes using the DS CLI

Complete this task to create count key data (CKD) volumes.

A logical volume consists of one or more data extents that are allocated from a single extent pool. The volume data type is inherited from the extent pool extent storage type (fixed block or CKD) characteristic. When a CKD volume is created, volume attributes are further defined by a base or alias volume type, 3390 or 3380 volume cylinder type, and volume capacity in cylinders. These volume attributes characterize the volume to the host system that will eventually access the volume. Each volume is assigned a volume ID, which is the volume address within the 64 KB address space. Host access to a volume is enabled when the volume ID is assigned to a volume group; however, CKD volumes are automatically assigned to the volume group CKD FICON/ESCON All (ID V10).

Complete the following steps to create your CKD volumes:

1. View your list of CKD extent pool IDs and determine which extent pool IDs that you want to use as the source for the CKD volumes to be created. You obtained this list when you first created your extent pools. If this list is not available, you can issue the **lsextpool** command to obtain the list of extent pool IDs.

Extent pool attributes determine the size and quantity of volumes that can be created. The extent pool ID (even/odd) indicates the storage server (0|1), which dictates that the logical control unit (LCU) ID component of the volume ID must be an even or an odd number.

2. Issue the **mkckdvol** command to make 128 base volumes for each LCU. Enter the **mkckdvol** command at the dscli command prompt with the following parameters and variables:

dscli> mkckdvol –dev IBM.2107-75FA120 -extpool p1 –cap 3339 —name finance#d 0000-007F

The following considerations affect the command example in this step:

- The **-extpool** parameter identifies a CKD extent pool that contains available data extents.
- The **-cap** parameter specifies the quantity of CKD cylinders that are allocated to this volume.
- The **-name** parameter allows you to assign an easy-to-use label or nickname to the volume. The volume name parameter can include a wildcard (#d or #h) that inserts a decimal or hexadecimal volume ID value into the volume name.

Note: The decimal designation does not apply to the volume ID number or the number of volumes that were created by the command. It only applies to the unique name that you have assigned to the volume. When you process the **mkckdvol** command, the volume name that you have assigned does not appear in the confirmation message. To view the volume name that you have assigned, issue the **lsckdvol** or **showckdvol** command.

- Volume ID 0000 007F specifies 128 volumes, starting at CKD address group (0), LCU ID (00), and volume number (00). You must specify volume IDs that have not been previously defined as CKD or fixed block volumes.
- **3**. Press Enter to create the volumes. A confirmation message is displayed that lists the successful creation of each volume.
- 4. Repeat Steps 2 and 3 until all required logical volumes for all LCUs have been created.

5. Issue the mkaliasvol command to make 128 alias volumes for each LCU. Enter the mkaliasvol command at the dscli command prompt with the following parameters and variables: dscli> mkaliasvol -dev IBM.2107-75FA120 -base 0000-004F -order decrement -qty 2 00FF

Consider the following conditions with regard to the command example in this step:

- The **-base** 0000 004F parameter specifies that alias volumes are assigned to existing base volume IDs 0000 004F. Base and alias volumes must be associated with a common LCU ID.
- The **-order** parameter specifies the order in which alias volume IDs are assigned.
- The **-qty** parameter specifies the number of alias volumes that are assigned to each base volume.
- The **volume ID** (00FF) parameter specifies that the alias volumes are assigned, starting at a CKD address group (0), LCU ID (00) and volume number (FF). You are responsible for specifying the volume ID values that have not been previously defined as CKD or fixed block volume types. As a result, alias volumes 00FF and 00FE are created for base volume 0000, 00FD and 00FC for 0001, and so on.
- 6. Repeat Step 5 until you have defined all required logical volumes for all the LCUs.
- 7. Press Enter to create the alias volumes. A confirmation message is displayed that lists the successful creation of each volume.
- 8. Issue the **lsckdvol** command to display a report that you can use to confirm the status of your CKD volumes. Enter the **lsckdvol** command at the dscli command prompt with the following parameters and variables:

dscli> lsckdvol -dev IBM.2107-1300861 -l 1410

Note: It is possible that the report will display that there was a configuration error that is associated with one or more of your **mkckdvol** transactions. In the majority of instances, the only way to correct this error is to issue the **rmckdvol** command.

Correcting a CKD volume configuration error

Complete this task to correct a count key data (CKD) volume configuration error.

There might be occasions when you are using the **mkckdvol** command to create CKD volumes, but the transaction fails. You might not be aware of the failure until you run the **lsckdvol** or the **showckdvol** command to check the status of the volumes that you have created.

The **lsckdvol** or the **showckdvol** commands display reports that includes a *configstate* category. The configuration state category reports on the current state of the volume. One of the configuration state codes is **configuration error**.

A status of **configuration error** specifies that the configuration process did not complete successfully. This state reflects an internal error condition and is not an indication that there was a user input error.

You might want to gather additional information about what caused the error, which can help you determine how to correct it. To correct this error state, you must delete the designated volume configuration and submit a new transaction request.

Complete the following steps to obtain additional information about the configuration error and to correct this error condition.

1. Add the **-v** (verbose) command parameter to your **mkckdvol** command and reissue the **mkckdvol** command for the transactions that show there is a configuration error.

Note: You can also turn on the verbose mode in your profile file and reissue the command. If you designate the verbose mode, the display of extra output includes the error code that is generated when the create CKD volume transaction fails.

2. Issue the **rmckdvol** command to delete the designated volume configurations if you do not want to obtain additional information about what caused the configuration error.

Note: In the majority of instances, this is the only method for correcting a configuration error.

Configuring Fibre Channel I/O ports using the DS CLI

Complete this task to configure Fibre Channel I/O ports using the DS command-line interface.

Before you begin, you must have the command-line interface prompt, and you must be connected to a storage image that is to be used for open systems host system storage.

In this process, you must designate the topology for the I/O port. The following three topology settings are available:

fibre channel arbitrated loop (coded as fc-al in the setioport command)

Enables the SCSI ULP with a FC-AL topology. The FC-AL topology does not support PPRC path I/O operations.

SCSI-FCP (coded as scsi-fcp in the setioport command)

Enables the SCSI ULP with a point-to-point or switched fabric topology. PPRC path I/O operations are enabled for this setting.

Note: Designate this topology for System i models using i5/OS level V5R3M5 and above.

ficon (coded as ficon in the setioport command)

Enables the FICON ULP with a point-to-point or switched fabric topology. PPRC path I/O operations are not supported for FICON ULP.

The storage image supports the Fibre Channel host adapter (HA) card type. For DS8000 machine types, HA cards are installed in I/O enclosures, each containing up to four HA cards. For DS6000 machine types, one or two HA cards are installed in each of the two processor node assemblies. Use the **lsioport** and **setioport** commands to configure Fibre Channel I/O ports.

Each Fibre Channel HA card contains four I/O ports. The storage image microcode automatically creates one I/O port to represent each HA card I/O port. The default Fibre Channel I/O port settings enable SCSI-FCP "identified" access to fixed block volumes. You might have to modify the I/O port settings to enable SCSI FC-AL access to fixed block volumes.

To configure the Fibre Channel ports, complete the following steps:

1. View a list of Fibre Channel port IDs by typing the following command format at the dscli command prompt:

dscli> lsioport -dev IBM.2107-75FA120 -l -type fc

A detailed report is displayed that lists the Fibre Channel I/O ports.

2. Analyze the report and determine which I/O port IDs that you want to access the fixed block volumes.

Configure a minimum of four I/O ports for SCSI host I/O operations. Select ports with physical locations on different host bus adapter (HA) cards. If possible, locate the HA cards in different I/O enclosures.

3. Set the I/O ports that you have identified to enable the FC-AL (Fibre Channel arbitrated loop), SCSI-FCP, or FICON topology. The following example shows how to enable the FC-AL topology by typing the following command format at the dscli command prompt:

Note: I/O ports are automatically set to the offline state and returned to the online state after configuration changes are applied.

dscli> setioport -dev IBM.2107-75FA120 -topology fc-al 0012 0013 0112 0113 4. Press Enter. A successful process returns a confirmation message indicating that the port IDs have been successfully configured.

Managing your logical storage configuration

After the initial creation of the configuration, there are additional tasks that must be processed to complete the configuration. Scan the following topics for the tasks that apply to your system.

Using DS CLI commands on i5/OS

Complete this task to use DS CLI commands from the "green screen" interface on i5/OS.

Before you can use the DS CLI on i5/OS, ensure that the following conditions have been met:

- You have installed the DS CLI code on i5/OS.
 - The DS CLI is installed on the i5/OS integrated file system (IFS) in the following two places:
 - IFS directory IBM/DSCLI, which contains the profiles, executable files, and readme files.
 - The library QDSCLI, which contains executable code.
- You have added library QDSCLI to the i5/OS library list by completing the following process:
 - 1. Enter WRKSYSVAL QUSRLIBL at the i5/OS command line.
 - 2. Press Enter and select option number 2.
 - 3. Add the QDSCLI library into the lib list.
- You have completed the initial configuration from the server from which you did your installation. For example, you have activated your licenses, created your arrays, ranks, extent pools, host attachments, and logical volumes on the DS8000 or DS6000.
- You have configured the DS CLI profile as appropriate. To edit the profile file, complete the following steps:
 - 1. Enter EDTF '/ibm/dscli/profile/dscli.profile' at the i5/OS command line.
 - 2. Update the following two lines when the profile file displays:
 - HMC IP address field and remove the comment from this command line
 - Dev ID field and remove the comment from this command line
- You have completed an IPL to System i5[®].

Note: Ownership and security on the QDSCLI library should be modified after installation to meet your security requirements.

Assuming that you have met the previous conditions, complete the following steps using DS CLI on i5/OS to process storage configuration and Copy Services functions on the DS8000 or DS6000:

1. From the i5/OS main menu, enter DSCLI at the prompt to start DS CLI on i5/OS and press Enter.

MAIN	OS/400 Main Men	u System:	IBMSYSTEM
Select one of the following:			
1. User tasks			
2. Office tasks			
3. General system tasks			
4. Files, libraries, and folders			
5. Programming			
6. Communications			
7. Define or change the system			
 8. Problem handling 9. Display a menu 			
10. Information Assistant options			
11. iSeries Access tasks			
90. Sign off			
Selection or command			
===> dscli			
F3=Exit F4=Prompt F9=Retrieve	F12=Cancel	F13=Information Assist	tant
F23=Set initial menu			

 The DSCLI displays the following screen where you can specify a DS CLI script for DS CLI commands and a DS CLI profile. In this example, a default profile is specified. The profile is not configured so the value of *DEFAULT is used. If you are not using a script, specify *None and press Enter.

After you press Enter, more fields appear in the screen as shown in step 3.

```
      Run DSCLI Functions (DSCLI)

      Type choices, press Enter.

      Script: *NONE or name . . . . *NONE____

      Profile . . . . . . . . . *DEFAULT

      F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display F24=More keys
```

3. Specify the values, as appropriate and press Enter.

```
Run DSCLI Functions (DSCLI)
Type choices, press Enter.
Script: *NONE or name . . . . > *NONE
Profile . . . . . . . . . . . .
                                *DEFAULT
                                *PROFILE
HMC1 . . . . . . . . . . . . . . .
HMC2 . . . . . . . . . . . . . . .
                                *PROFILE
                                admin
Password . . . . . . . . . . . .
Install Path . . . . . . . . . .
                                '/ibm/dscli'
DSCLI CMD . . . . . . . . . . .
                                *int
                                                                         Bottom
F3=Exit F4=Prompt F5=Refresh F12=Cancel F13=How to use this display
F24=More keys
```

Consider the following fields and values:

- If you are using a DS CLI script for DS CLI commands, enter the name in the Script field. Otherwise, specify ***None**.
- If you use default profile, leave the value *DEFAULT in the field Profile. If you use another file as a profile, specify the name and path of this file in the field Profile.
- Enter the hardware management console (also known as the management console) user in the User field. Typically, it is Admin.

- Enter the password of the user (typically the administrator's password).
- Enter *INT (for interactive session) in the DSCLI CMD field.
- The screen shown in step 4 displays.
- 4. Specify the DS CLI commands to invoke storage configuration or Copy Services functions.

```
dscli>
===>
F3=Exit F6=Print F9=Retrieve F12=Exit F13=Clear F17=Top F18=Bottom
F21=CL command entry
```

Modifying an extent pool

Complete this task to modify the properties of an extent pool using the DS CLI.

Use the **chextpool** command to modify the properties of an extent pool. You can modify the following extent pool properties:

- name of the extent pool
- extent limit on or off indicator
- extent limit percentage
- extent threshold percentage

Complete the following steps to modify the extent pool properties.

1. Issue the **lsextpool** command to generate a report that identifies the status of your extent pools by storage type (fixed block or count key data). Enter the **lsextpool** command at the dscli command prompt with the following parameters and variables:

dscli> lsextpool -dev storage_image_ID -l -stgtype (fb | ckd)

Note: The **-stgtype** parameter must be designated as either *fb* (fixed block) or *ckd* (count key data). The storage type allows you to limit the list of extent pools for issues such as which ones to rename or to change the limit or threshold percentages.

2. Issue the **chextpool** command to change the name that is assigned to the extent pool or to change the percentages that are allocated for extent and threshold limits. Enter the **chextpool** command at the dscli command prompt with the following parameters and variables:

```
dscli> chextpool -dev storage_image_ID -name new_extent_pool_name
-extentlimit [on | off] -limit extent_limit_percentage
-threshold extent threshold percentage -extentpool ID
```

Notes:

- a. The new extent pool name can include up to sixteen characters.
- b. The **-extentpool_ID** parameter is required but does not need to be specified as a separate entry. You can add it to the *storage_image_ID* parameter. For example: IBM.2107-75FA120/P21, with P21 being the extent pool ID. Extent pool IDs are specified as 4-digit values with no leading zeros, and they are preceded by the letter P.
- c. The unique name that you assigned to the extent pool does not display in the output message of the chextpool command. However, when you issue the lsextpool command, the extent pool name is displayed.
- 3. Issue the **lsextpool** command to verify that your changes have been processed.

Viewing extent pool status

Complete this task to display a list of extent pools in a storage system and status information on each extent pool in the list.

Enter the **lsextpool** command if you want to view the unique names that you have assigned to your extent pools or if you want to view general status information about the extent pools in your storage system. If you want to view the details or properties that are associated with your extent pools or if you want to view the performance metrics, enter the **showextpool** command.

Complete the following step to display a list of extent pools and their status in a storage system.

Enter the **lsextpool** command to display the extent pool list and status information. dscli> lsextpool -dev *storage_image_ID* -l

Notes:

- 1. Use the **-1** parameter if you want to see the list and status for all the extent pools (fixed block and CKD) in your storage system. A full report is displayed.
- 2. Use the **-s** parameter if you want to see only a list of the extent pools in your storage system. No additional information is provided.

Viewing extent pool properties and performance metrics

Complete this task to display the detailed properties for the list of extent pools in a storage unit and to view the performance metrics status information on each extent pool in the list.

You must know and use an extent pool ID that resides in your storage system. You can obtain these IDs by entering the **lsextpool** command.

Enter the **showextpool** command when you want to view the details of the properties that are associated with an extent pool or when you want to view the performance metrics for an extent pool in your storage system.

Complete the following steps to display the detailed properties of an extent pool or to display the performance metrics of an extent pool.

- (For detailed properties information) Enter the showextpool command. dscli> showextpool -dev storage_image_ID extentpool_ID
- 2. (For performance metrics information) Enter the showextpool command. dscli> showextpool -dev storage_image_ID -metrics extentpool_ID

Notes:

- a. All performance metrics are an accumulation since the most recent counter wrap or counter reset.
- b. The extent pool performance counters are reset on the following occurrences:
 - When the storage system is turned on.
 - When a server has failed, and the failover and failback sequence is performed.

Deleting extent pools from a storage configuration

Complete this task to remove one or more extent pools from a storage configuration.

When you are using the DS CLI to delete extent pools as part of a storage configuration deletion, the following sequential deletions must have already occurred:

- The volumes associated with the extent pool must be removed.
- (CKD volume configuration only) The logical control units (LCUs) that are associated with the extent pool must be removed.
- The ranks that are assigned to the extent pool must be unassigned or removed.
- The arrays that are assigned to the extent pool must be removed.

To delete an extent pool or a number of extent pools, you must first generate a list of extent pool IDs by storage type (fixed block or CKD) by issuing the **lsextpool** command. After you determine which extent pools can be deleted, you can issue the **rmextpool** command that designates the extent pools that you want to delete.

Complete the following steps to delete one or more extent pools from a fixed block volume configuration.

1. Issue the **lsextpool** command to display a list of extent pools. Ensure that you designate the storage type within your command parameters. Enter the **lsextpool** command at the dscli command prompt with the following parameters and variables:

dscli> lsextpool -dev storage_image_ID -l -stgtype [fb | ckd]

- 2. Analyze the list and determine which extent pools can be deleted.
- **3.** Issue the **rmextpool** command to delete the designated extent pools. Enter the **rmextpool** command at the dscli command prompt with the following parameters and variables:

dscli> rmextpool -dev storage_image_ID extentpool_ID

Note: If you are deleting several extent pools, you can add the **-quiet** parameter to your command. This parameter turns off the confirmation message that is generated for each deletion transaction.

4. Issue the **lsextpool** command after the deletion processing has completed to verify that the extent pools have been deleted.

Viewing the array disk drive module status

Complete this task to view the status of an array disk drive module (DDM) using DS CLI commands.

The DS8000 machine type storage image DDMs are packaged into storage enclosure pairs. The DS8000 machine type contains at least one storage enclosure pair, with a minimum of 16 DDMs. All DDMs that are installed in a storage enclosure pair have identical capacity, rpm (revolutions per minute), and interface characteristics.

The DS6000 machine type contains at least one storage enclosure, with a minimum of four DDMs.

The DDMs of a storage enclosure are partitioned into array sites. A DS8000 machine type array site consists of eight DDMs, four from each storage enclosure of a storage enclosure pair, two-or-four (eight DDM) array sites per storage enclosure pair. A DS6000 machine type array site consists of four DDMs in one storage enclosure of a storage enclosure pair, with two-to-eight (four DDM) array sites per storage enclosure pair. All array sites of a storage enclosure pair have identical capacity, rpm, and interface characteristics, and an interface to a common DA pair.

The DDMs of a storage enclosure are partitioned into array sites. The creation of arrays is based on the array sites that are associated with the storage unit. Before and after creating an array, you might want to check on the status of the DDMs.

Complete the following steps to view the DDM status:

1. Issue the **1sddm** command to obtain a list and status of the DDMs currently associated with the storage unit. Enter the **1sddm** command at the dscli command prompt with the following parameters and variables:

dscli> lsddm -l storage_image_ID

 Issue the showarraysite command after you have created an array using the DDMs. Enter the showarraysite command at the dscli command prompt with the following parameters and variables: dscli> showarraysite storage_image_ID -fullid site_ID

Notes:

- **a**. The storage image ID is optional. You do not have to specify it but, if you choose not to use it, you must provide a fully qualified *site_ID* which includes the manufacture, model type, and serial number information.
- b. The *site-ID* parameter is a four-digit number preceded by the letter "S" with no leading zeros.
- **c**. The **showarraysite** command provides the following DDM information that is associated with the DDM after the array has been created:
 - DDM serial number
 - Spares Identifies, if any, the number of spare DDMs that are allocated from the array site.
 - Data DDM Specifies the number of data DDMs. This value is based on the number of DDMs minus the number of spares.

Viewing array status

Complete this task to view the status of all arrays that are associated with a storage unit using DS CLI commands.

The steps in this task presume that you have already created your arrays from the array sites.

The creation of arrays is based on the array sites that are associated with the storage system. After you have created your arrays, there might be times when you want to view the status of the array sites and the associated arrays.

Complete the following steps to view the status of all the array sites and arrays that are associated with the storage system.

1. Enter the **lsarraysite** command to generate a list of all the array sites and their status:

dscli> lsarraysite -dev storage_image_ID -l

The state column of the report might be of interest as it specifies the following state of the array site and conditions that require attention:

For DS8000:

- 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 are in the Normal state.
- Unavailable The array is not assigned to a rank and one or more of the disk drive modules are not in the Normal state.

For DS6000:

- Assigned The array site has been defined as an array.
- Unassigned 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.
- 2. Issue the **lsarray** command to generate a list of all the arrays and their status. Enter the **lsarray** command at the dscli command prompt with the following parameters and variables: dscli> lsarray -dev *storage image ID* -1

Note: You might want to analyze the state and data column information for the arrays. Some of the reported conditions require further action. See the **lsarray** command for additional information.

Viewing properties for one array

Complete this task to view the detailed properties of one array and the array site that is associated with a storage system by using DS CLI commands.

The steps in this task presume that you previously created your arrays.

The creation of arrays is based on the array sites that are associated with the storage system. After you create your arrays, there might be times when you want to view the status of an array site and the associated array.

Note: For your DS6000 machine type, you might create your array from one or two array sites. An array inherits the characteristics of its parent array sites and is given a RAID type attribute (5 or 10).

Complete the following steps to view the status of the array site or sites and the array that is associated with the storage system.

1. Enter the **showarraysite** command to generate a report that displays the array site or sites and their status.

dscli> showarraysite -dev storage_image_ID site_ID

Notes:

- a. The site ID is a four-digit number that is preceded by the letter *S* with no leading zeros.
- b. The site ID does not specify a physical location. It is, however, an identifier for the array site ID.
- c. If you created the array previously, the array site state shows a value of assigned.
- 2. Enter the **showarray** command to generate the properties report for the specified array. dscli> showarray -dev *storage_image_ID array_ID*

Removing arrays from a storage configuration or a rank assignment

Complete this task to remove an array or a range of arrays from a storage configuration using the DS CLI.

When the array or arrays are deleted as part of deleting a storage configuration and you are using DS CLI to delete the configuration, the following sequential deletions must have already occurred:

- Host access to the volumes of the configuration must have been removed. (Does not apply to a CKD configuration.)
- The associated volume groups must have been removed. (Does not apply to a CKD configuration.)
- The fixed block or CKD volumes that are part of the configuration must have been removed.
- The LCUs (if you are removing a CKD configuration) must have been removed.
- The ranks must have been removed.

There might be times when you want to remove arrays from a storage configuration or from a rank assignment. You can avoid errors when you use the DS CLI by ensuring that the arrays are ready for removal. The arrays must have a status of "unassigned" before they can be removed or be reassigned to another rank.

When you remove a rank using the DS CLI, there is an extended period of processing because the array is unassigned from the rank and the drives are formatted. During this processing, the array is still shown with a status of assigned, even though you have received a confirmation message that the rank has been removed. The status for the array does not change to "unassigned" until after the array has been formatted.

Complete the following steps to remove arrays from a storage configuration or a rank:

1. Issue the **lsarray** command to obtain a list of array IDs to be removed. Enter the **lsarray** command at the dscli command prompt with the following parameters and variables:

dscli> lsarray -dev storage_image_ID -state unassigned

Notes:

- a. You might have to issue the **lsarray** command several times before you observe that the arrays are in a state that allows them to be removed or reassigned.
- b. Specify the **-state***unassigned* parameter to narrow your list to just the array IDs that are not assigned to a rank ID.
- c. If you issue the **lsarray** command without using the **-state***unassigned* parameter, you might see a list of arrays that have a state of unavailable. This is a good indication that the ranks have not been removed and that the arrays are still formatting. You must wait until the ranks have been removed and the arrays have been formatted before you can proceed.
- d. Proceed to the next step (remove arrays) only after all the arrays that you want to remove or reassign are displayed with a state of unassigned.
- 2. Issue the **rmarray** command to delete the unassigned arrays so that the array sites can be redefined as new arrays. Enter the **rmarray** command at the dscli command prompt with the following parameters and variables:

dscli> rmarray -dev storage_image_ID array_ID

Notes:

- a. You can remove one or many arrays as long as you designate the range of arrays using a hyphen and separate each range of arrays or a single array with a space before the next array designation. For example, A44-A48 A51 designates a range of arrays and a single array.
- b. If you are removing several arrays, you might want to designate the **-quiet** parameter in your command. This parameter turns off the deletion confirmation message that is generated after each array is deleted.

Adding a rank to an extent pool

Complete this task to add an unassigned rank to an extent pool using the DS CLI.

To add a rank to an extent pool, the rank must have a data state designation of normal and a configuration state designation of unassigned.

A rank is a logically contiguous storage space made up of one or more arrays. An unassigned rank is not associated with either rank group 0 or 1. Any unassigned rank can be assigned to an extent pool that is associated with either rank group 0 or 1. Over time, you might remove a rank from an array and extent pool without deleting the rank. When a rank is removed and not deleted, it retains its storage type designation of fixed block or CKD. This designation cannot be changed.

Complete the following steps to add a rank to an extent pool:

1. Issue the **1srank** command to generate a report that lists the status of the ranks that are associated with the storage unit. Enter the **1srank** command at the dscli command prompt with the following parameters and variables:

dscli> lsrank -dev storage_unit_ID -l -state unassigned

Notes:

- **a**. The report that is generated by this example provides a list of all unassigned ranks; however, the storage type is mixed between fixed block and CKD.
- b. You can narrow your report information to a specific storage type by adding the **-stgtype** [fb | ckd] parameter to your command.
- 2. Issue the chrank command to add (reassign) a rank or ranks to an extent pool. Enter the chrank command at the dscli command prompt with the following parameters and variables: dscli> chrank -dev storage_image_ID -extpool extentpool_ID rank_ID

Notes:

- **a**. The rank ID is a four-digit number with the prefix *R* and with no leading zeros. You can specify a range of rank IDs by using a hyphen between the beginning and ending values of the range. For example: R102-R105
- b. You can specify multiple rank IDs or rank ID ranges, but you must leave a space between each designation. For example: R102 R105 R107-R109

Modifying a rank

Complete this task to modify a rank using the DS CLI.

You can complete the following modifications to a rank using the DS CLI commands:

- Designate that the rank be given a reserved status.
- Release a rank from a reserved status.
- Designate that the rank be removed (but not deleted) from its current extent pool and array assignment and be designated as unassigned.
- Designate that the rank be assigned to an extent pool.

Complete the following steps to modify a rank using DS CLI commands:

1. Issue the **1srank** command to generate a report that lists the status of the ranks that are associated with the storage unit. Enter the **1srank** command at the dscli command prompt with the following parameters and variables:

dscli> lsrank -dev storage_unit_ID -l

Notes:

- a. The report that is generated by this example provides a list of all ranks; however, the storage type is mixed between fixed block and CKD.
- b. You can narrow your report information to a specific storage type by adding the **-stgtype** [fb | ckd] parameter to your command.
- 2. Use the report to determine the rank or ranks you want to modify. The report contains details about the ranks that you must use to issue the **chrank** command for modifications.
- 3. Issue the **chrank** command to implement one of the following types of modifications:
 - **a**. To designate a rank as reserved, enter the **chrank** command at the dscli command prompt with the following parameters and variables:

dscli> chrank -dev storage_image_ID -reserve rank_ID

Changing the rank configuration state to "reserved" designates that the extents that are associated with the rank are not eligible for allocation to a logical volume. However, the existing allocations remain in effect until the configuration state is changed to normal (the characteristics that the rank inherited from its parent array when it was originally assigned remain intact).

Notes:

- 1) You can specify a range of rank IDs or multiple rank IDs as long as you match the command usage criteria.
- 2) You cannot change the configuration state of a reserved rank to "unassigned" without first releasing it.
- **b.** To release a rank from its reserved configuration state, enter the **chrank** command at the dscli command prompt with the following parameters and variables:

dscli> chrank -dev storage_image_ID -release rank_ID

When a rank is released from the configuration state of "reserved", it is designated with a configuration state of "normal".

c. To remove a rank from its current extent pool and array assignment but not delete it, enter the chrank command at the dscli command prompt with the following parameters and variables: dscli> chrank -dev storage image ID -unassign rank ID

Notes:

- 1) A rank must have a configuration state of *normal* before it can be changed to a configuration state of *unassigned*.
- 2) A rank that is unassigned can be assigned to an array and extent pool of another storage configuration as long as the storage type is compatible: all fixed block or all CKD.

Viewing rank status

Complete this task to view the status of all the ranks that are associated with a storage system using DS CLI commands.

After you have created your ranks, there might be times when you want to view the status of all the ranks that are associated with your storage system. Using the parameters that are associated with the **lsrank** command, you can refine your search to specific rank criteria such as:

- Storage type (fixed block or CKD)
- Data state
- Configuration state
- RAID type

Complete the following step to view the status of all the ranks that are associated with the storage system.

Enter the **lsrank** command to generate a list of all the ranks and their status. dscli> lsrank -dev *storage image ID* -1

The state and datastate column information for the ranks contains reported conditions that can require further action. See the **lsrank** command for an explanation of the action designations.

Viewing properties for one rank

Complete this task to view the detailed properties of one rank that is associated with a storage system using DS CLI commands.

A rank is a logically contiguous storage space that is made up of one array. After you have created your ranks, there might be times when you want to view the status of an individual rank.

Complete the following steps to view the status of a rank that is associated with the storage system.

Enter the **showrank** command to generate the properties report for the specified rank. dscli> showrank -dev *storage_image_ID* rank_ID

Notes:

- 1. Because the **showrank** command requires the use of a specific rank ID, you can enter the **lsrank** command first to obtain the specific rank IDs.
- 2. The state and datastate column information for the ranks contains reported conditions that can require further action. See the **showrank** command for an explanation of the action designations.

Correcting a rank-related configuration error

Complete this task to correct a rank-related configuration error.

There might be occasions when you are using the **mkrank** command to create ranks, but the transaction fails. You might not be aware of the failure until you run the **lsrank** or **showrank** command to check the status of the ranks that you have created.

The **lsrank** or the **showrank** commands display reports that includes a *state* category. The state category reports on the current state of the rank. One of the state codes is **configuration error**.

A state of **configuration error** specifies that a rank configuration process has not completed successfully. This state reflects an internal error condition and is not an indication that there was a user input error.

You might want to gather additional information about what caused the error, which can help you determine how to correct it. To correct this error state, you must delete the designated rank configuration and submit a new transaction request.

Complete the following steps to obtain additional information about the configuration error and to correct this error condition.

- 1. Obtain additional information about the transaction by implementing one of the following methods:
 - Add the **-v** (verbose) command parameter to your **mkrank** command and reissue the command for the transactions that show the configuration error designation.

Note: You can also turn on the verbose mode in your profile file and reissue the command. Designating the verbose mode allows the display of extra output that includes the error code that is generated when the create rank transaction fails.

• Add the **-extpool** parameter to your **mkrank** command and reissue the command for the transactions that show the configuration error.

You might consider using this parameter if you have not yet assigned your ranks to the extent pools. If the transaction fails, a message states the reason for a failure.

2. Issue the **rmrank** command to delete the designated rank configurations if you do not want to obtain additional information about what caused the configuration error.

Note: In the majority of instances, this is the only method for correcting a configuration error.

Removing ranks from a storage configuration

Complete this task to remove ranks from a storage configuration using the DS CLI.

When you are using the DS CLI to delete ranks as part of a storage configuration deletion, the following sequential deletions must have already occurred:

- Host access to the volumes of the configuration must have been removed. (Does not apply to a CKD configuration.)
- The associated volume groups must have been removed. (Does not apply to a CKD configuration.)
- The fixed block or CKD volumes that are part of the configuration must have been removed.
- The LCUs (if you are removing a CKD configuration) must have been removed.

When you remove ranks using the DS CLI, there is an extended period of processing because the arrays and extent pools are unassigned from the ranks and the drives are formatted. During this processing, the arrays and extent pools are still shown with a status of assigned, even though you receive a confirmation message each time a rank has been deleted. The status for the arrays and extent pools do not change to "unassigned" until after the drives have been formatted.

Complete the following steps to remove ranks from a storage configuration:

1. Issue the **1srank** command to obtain a list of the ranks that are associated with the storage configuration that is being deleted. Enter the **1srank** command at the dscli command prompt with the following parameters and variables:

```
dscli> lsrank -dev storage_image_ID -l -stgtype [fb | ckd]
```

2. Look at the list and ensure that the ranks are in a state that allows them to be deleted. All the ranks need to have a data and configuration state of normal.

3. Issue the **rmrank** command to delete the ranks from the storage configuration. Enter the **rmrank** command at the dscli command prompt with the following parameters and variables: dscli> rmrank -dev storage image ID rank ID

Notes:

- a. If you have multiple ranks that are being deleted, you might want to include the **-quiet** parameter in your command. This parameter suppresses the confirmation message that is issued for each rank that is deleted.
- b. Deleting a rank or many ranks is a lengthy process because the array and extent pool assignments are unassigned and the disk drives are formatted.

When a rank is unassigned from the array and extent pool, a confirmation messages is issued that indicates that the rank has been deleted. However, because of the formatting, the process is not complete. You cannot initiate any action on the arrays or extent pools until the formatting is completed.

Modifying a logical control unit

Complete this task to modify a logical control unit (LCU) using the DS CLI.

You can complete the following modifications to an LCU using the DS CLI commands:

- Change the subsystem ID to ensure it retains its unique identity
- Change the system behavior so that it emulates an LCU type that allows your system to process DS CLI transactions
- Change the behavior of consistency group creation
- · Change the system behavior for processing concurrent copy transactions
- Change the system behavior for processing extended remote copy transactions (for DS8000)

Complete the following steps to modify an LCU using DS CLI commands:

1. Issue the **lslcu** command to generate a report that lists the status of the LCUs that are associated with the storage unit. Enter the **lslcu** command at the dscli command prompt with the following parameters and variables:

dscli> lslcu -dev storage_unit_ID -l

- 2. Use the report to determine the LCU or LCUs that you want to modify. The report contains details about the LCUs that you must use to issue the **chlcu** command for modifications.
- 3. Issue the **chlcu** command to implement one of the following types of modifications:
 - a. To maintain the unique identity that is associated with your logical subsystem within your Copy Services domain, you can change your subsystem ID (SSID). Enter the **chlcu** command at the dscli command prompt with the following parameters and variables:

```
dscli> chlcu -dev storage_image_ID -ss new_ss_ID lcu_ID
```

Note: The new SSID that you specify replaces the existing SSID value in the initial target LCU ID.

b. To provide your system a format that allows you to process DS CLI transactions. Enter the chlcu command at the dscli command prompt with the following parameters and variables:
 dscli> chlcu -dev storage_image_ID -lcutype [3990-3 | 3990-tpf | 3990-6 | bs2000] lcu ID

Notes:

- 1) The target LCUs are changed to the LCU type that you designate.
- 2) When you designate multiple LCUs, separate multiple IDs and multiple ID ranges with a space. Separate your LCU range with a dash (-) between the first and last number of the range.
- **c**. To modify the concurrent copy timeout value using the **chlcu** command, see "Modifying the Concurrent Copy timeout value" on page 574.

- d. To modify the consistency group timeout value, see "Modifying the consistency group timeout value" on page 574.
- e. To modify the critical mode (administrator authority only), see "Modifying the critical mode setting" on page 575.
- f. To modify the z/OS Global Mirror timeout value for DS8000, see "Modifying the z/OS Global Mirror timeout value" on page 575.

Viewing logical control unit status

Complete this task to view the status of all the logical control units (LCUs) that are associated with a storage system using DS CLI commands.

After you have created your LCUs, there might be times when you want to view the status of all the LCUs that are associated with your storage system. Using the parameters that are associated with the **lslcu** command, you can refine your search to the following specific LCU criteria:

- Address group
- Specific LCUs or multiple LCUs

Complete the following step to view the status of all the LCUs that are associated with the storage system.

Enter the **lslcu** command to generate a list of all the LCUs and their status.

dscli> lslcu -dev *storage_image_ID* -l

Notes:

- 1. Enter the **lsaddressgrp** command first if you decide to refine your search to include just the LCUs that are associated with a specific address group. The **lsaddressgrp** command provides a list of address groups that you can then use with the **-addrgrp** parameter of the **lslcu** command.
- **2**. To specify a range of LCU IDs, separate the LCU IDs with a dash (-). You must separate multiple LCU IDs or ranges of LCU IDs with a blank space between each ID or range of IDs.

Viewing properties for one logical control unit

Complete this task to view the detailed properties of one logical control unit (LCU) that is associated with a storage system using DS CLI commands.

An LCU represents a logical subsystem for z Systems hosts.

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.

The DS6000 has a 16 384 volume address space that is partitioned into 64 logical subsystem (LSS) units, where each LSS contains 256 logical volume numbers. The 64 LSS units are assigned to one of 4 address groups, where each address group contains 16 LSSs, or 4096 volume addresses. All of the LSSs in one address group must be of the same type (CKD or fixed block).

Because you can modify some of the properties of an LCU, there might be times when you want to examine the associated properties. The **showlcu** command allows you to view the properties of a single LCU.

Complete the following step to view the properties of a single LCU:

Enter the **showlcu** command to view a report that displays the properties of a single LCU. dscli> showlcu -dev *storage_image_ID LCU_ID*

Note: The LCU ID is a two-digit hexadecimal number in the range of 00 - FE for DS8000 or 00 - 1F for DS6000.

Removing logical control units from a CKD storage configuration

Complete this task to remove all logical control units (LCUs) from a CKD storage configuration using the DS CLI.

When you are using the DS CLI to delete LCUs as part of a storage configuration deletion, the following sequential deletions must have occurred:

- The alias CKD volumes that are part of the configuration must have been removed
- The CKD volumes that are part of the configuration must have been removed

Complete the following steps to remove LCUs from a CKD storage configuration:

 Issue the lslcu command to obtain a list of the LCUs that are associated with the storage configuration that is being deleted. Enter the lslcu command at the dscli command prompt with the following parameters and variables:

dscli> lslcu -dev storage_image_ID -l

- 2. Look at the list to ensure that the LCUs are in a state to be removed. They are ready if there are no volumes that are assigned to the LCU (zeros are displayed for each LCU in the Confgvols column of the list).
- Issue the rmlcu command to delete the LCUs from the storage configuration. Enter the rmlcu command at the dscli command prompt with the following parameters and variables: dscli> rmlcu -dev storage_image_ID lcu_ID

Notes:

- a. If you have multiple LCUs that are being deleted, you can include the **-quiet** parameter in your command. This parameter suppresses the confirmation message that is issued for each LCU that is deleted.
- b. You must separate multiple LCU IDs or ranges of LCU IDs with a blank space between each ID or range of IDs. Each range of LCU IDs must be separated by a dash (-) between the first ID and the last ID of the range.

Chapter 7. Copy Services functions

This topic provides information about how to use DS CLI commands to complete the Copy Services tasks associated with FlashCopy, Metro Mirror, Path establishment, and Global Mirror transactions.

FlashCopy functions

This topic provides a list of tasks that help you create, monitor, and manage your FlashCopy operations using DS CLI commands.

Creating a FlashCopy relationship

Complete this task to create a FlashCopy relationship between a source and target volume that enables a point-in-time copy of a source volume onto a target volume.

You can create a FlashCopy relationship between a source and a target volume that enables a point-in-time copy of a source volume onto a target volume. FlashCopy functions run on a DS8000 model or DS6000 model and are supported on many operating systems. For example, if you set up and configure your machine to use i5/OS, you can create copies of System i disk pools within a single machine using FlashCopy. After the FlashCopy function completes, you can immediately access the target point-in-time copies by associating another System i or logical partition.

When you issue a FlashCopy command with the background copy option, the FlashCopy relationship is established but it is put in a queue for background copying. The time difference between the submission and actual start time of the task depends on the number of FlashCopy relationships that are currently copying in the background or waiting in the queue. When the copy processing starts, the status displays as "background copy running" for that FlashCopy volume pair.

How long the actual physical copy processing takes can depend on the amount of data being copied and other activities that are occurring on the storage unit.

Note:

- At the data set level, the maximum number of FlashCopy relationships allowed on a volume is 65534. If that number is exceeded, the FlashCopy operation fails.
- With full volume FlashCopy, a source volume can be in FlashCopy and FlashCopy Space Efficient relationships with a maximum of 12 target volumes. If that number is exceeded, the FlashCopy operation fails. Of the up to 12 FlashCopy relationships, only one incremental FlashCopy relationship is allowed.

With FlashCopy Space Efficient, a source volume can be in FlashCopy and FlashCopy Space Efficient relationships with a maximum of 12 target volumes. If that number is exceeded, the FlashCopy operation fails.

You can complete this task using either the DS CLI or the DS8000 Storage Management GUI.

Complete the following steps to create FlashCopy relationships using the DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information the command requires. The second format is an example command with declared values for the variables.

1. Enter the **mkflash** command to create FlashCopy relationships.

dscli> mkflash -dev storage_image_ID sourcevolumeID:targetvolumeID

Example

dscli> mkflash -dev IBM.2107-75FA150 0001:0004

Notes:

- **a**. Specify the storage unit for the **-dev** *storage_image_ID* parameter. This parameter is required if you do not specify a fully qualified ID for the source and target volumes and you do not specify a value for the devid variable in your profile file. If the management console has an IP connection to the specified storage unit, the command works. If the IP connection is not established, you can use the **mkremoteflash** command if there is a PPRC Path established between the storage unit from which you issue the command and the (remote) storage unit where the FlashCopy volumes are located.
- b. For further information, including optional parameters, see the **mkflash** and **mkremoteflash** commands.

A confirmation message is issued for each successful FlashCopy pair that is created.

2. Enter the lsflash command to check the status information for each FlashCopy relationship. A detailed report (when you use the -l parameter) is displayed for each FlashCopy relationship. dscli> lsflash -dev storage_image_ID -l SourceVolumeID:TargetVolumeID. dscli> lsflash -dev IBM.2107-75FA150 -l 0100:0100

Use the-1 parameter to provide a more detailed report about the FlashCopy relationships.

Note: If you used the **mkremoteflash** command, you must enter the **lsremoteflash** command to run a status check.

Creating a persistent FlashCopy relationship

Complete this task to create a persistent FlashCopy relationship that remains even after the FlashCopy operation completes.

Creating a persistent FlashCopy relationship prevents another FlashCopy task from writing to your target volume before you have deleted the FlashCopy relationship.

You can complete this task using either the DS CLI or the DS8000 Storage Management GUI.

Complete the following step to create a persistent FlashCopy relationship. The example commands in this task are shown in two formats. The first format shows the type of information the command requires. The second format is an example command with declared values for the variables.

Issue the **mkflash** command with the **-persist** parameter to create a persistent FlashCopy relationship. Enter the **mkflash** command at the dscli command prompt with the following parameters and variables:

dscli> mkflash -dev storage_image_ID -persist sourcevolumeID:targetvolumeID

Example dscli> mkflash -dev IBM.2107-75FA120 -persist 0100:0200

The resulting output

FlashCopy pair 0100:0200 successfully created.

Viewing information about FlashCopy relationships

Complete this task to view status information about each existing FlashCopy relationship.

You can complete this task using either the DS CLI or the DS8000 Storage Management GUI.

Complete the following step to view status information about existing FlashCopy relationships using DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format is an example command with declared values for the variables.

Issue the **lsflash** command to provide a report that lists the FlashCopy relationships and status information for each FlashCopy relationship in the list. Enter the **lsflash** command at the dscli command prompt with the following parameters and variables:

dscli> lsflash -dev storage_image_ID -l source_volume_ID:target_volume_ID

Example

dscli> lsflash -dev IBM.2107-75FA150 -l 0100:0200 0101:0201 0102:0202 0103:0203

The resulting output

Note: The following tables display the output that is associated with the **lsflash** command using the **-l** parameter.

ID	SrcLSS	Seq- uence Num	Time- out (secs)	Active Copy	Record ing	Persist ent	Rever tible
0100:0200	01	10	120	Disabled	Disabled	Disabled	Disabled
0101:0201	01	10	120	Disabled	Disabled	Disabled	Disabled
0102:0202	01	11	120	Disabled	Disabled	Disabled	Disabled
0103:0203	01	11	120	Disabled	Disabled	Disabled	Disabled

Source- Write- Enabled	Target- Write- Enabled	Back- ground- Copy	Copy- Indic- ator	OutOf- Sync- Tracks	Date- Created	Date- Synced
Enabled	Disabled	Disabled	Yes	0	12/01 /2003 02:20 :00	12/01 /2003 02:23 :47
Enabled	Disabled	Disabled	Yes	0	12/01 /2003 02:20:00	12/01 /2003 02:23:47
Enabled	Disabled	Disabled	Yes	0	12/01 /2003 02:20 :00	12/01 /2003 02:23 :47
Enabled	Disabled	Disabled	Yes	0	12/01 /2003 02:20 :00	12/01 /2003 02:23 :47

Deleting FlashCopy relationships

Complete this task to delete FlashCopy relationships.

Deleting FlashCopy relationships between volume pairs ends a FlashCopy operation. You can delete a FlashCopy relationship at any time. If you delete a FlashCopy relationship with the background copy option and the background copy operation is still in progress, the target volume is not a complete point-in-time copy of the source volume.

You can complete this task using either the DS CLI or the DS8000 Storage Management GUI.

Complete the following steps to remove FlashCopy relationships using DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information the command requires. The second format is an example command with values declared for the variables.

Issue the lsflash command to check the status information for each FlashCopy relationship. A
detailed report (when you use the -l parameter) is displayed for each FlashCopy relationship. Enter
the lsflash command at the dscli command prompt with the following parameters and variables:

dscli> lsflash -dev storage_image_ID -l SourceVolumeID:TargetVolumeID.

dscli> lsflash -dev IBM.2107-75FA150 -l 0100:0100

Use the -1 parameter to provide a more detailed report about the FlashCopy relationships.

Note: If you have originally used the **mkremoteflash** command to create your FlashCopy relationships, you must enter the **lsremoteflash** command to run a status check.

- **2**. Analyze the list of volumes and ensure that these are the volumes from which the FlashCopy relationship must be removed.
- **3**. Issue the **rmflash** command to remove the FlashCopy volume relationships. Enter the **rmflash** command at the dscli command prompt with the following parameters and variables:

dscli> rmflash -dev storage_image_ID SourceVolumeID:TargetVolumeID

```
Example
```

dscli> rmflash -dev IBM.2107-75FA120 0001:0004 0003:00FF 0008:000C

Notes:

- **a**. The example shows the use of multiple FlashCopy pair IDs. Ensure that you separate multiple FlashCopy pair IDs with spaces.
- b. If you used the **mkremoteflash** command to create your FlashCopy relationships, you must enter the **rmremoteflash** command to remove the FlashCopy relationships.
- 4. A confirmation message is displayed for each FlashCopy relationship that you want to remove. Enter **Y** in response to each message that requests that you confirm that you want to remove the specified FlashCopy pair. A message like the following example appears for each FlashCopy pair being removed when you process the **rmflash** command.

Are you sure you want to remove the FlashCopy pair 0001:0004? [y/n]: Y

FlashCopy pair 0001:0004 successfully removed.

Creating remote FlashCopy transactions

Complete this task to create a remote FlashCopy (inband FlashCopy on the ESS 2105) at a target (remote) site using remote FlashCopy commands.

Remote FlashCopy operations can only be processed using the DS CLI and not the DS8000 Storage Management GUI. (Part of the Remote FlashCopy operation requires that you create paths and volume pairs first. You can issue those requests using either the DS8000 Storage Management GUI or the DS CLI.)

To establish a FlashCopy relationship at the target site, remote FlashCopy commands are issued to a source volume of a remote mirror and copy volume pair on a source (local) storage unit and sent across paths (acting as a conduit) to a target storage unit. This eliminates the need for a network connection to the target site solely for the management of FlashCopy relationships.

Limitation: Remote FlashCopy commands establish a FlashCopy relationship at the target (remote) site when a network connection to the target site is lost. The Remote FlashCopy operation is not supported through the DS8000 Storage Management GUI, because network connections to both the source and target sites are required. If the network connection to the target site is lost, the DS8000 Storage Management GUI cannot connect to the target site. Whether you use the DS8000 Storage Management GUI or DS CLI for Steps 1 and 2, you must perform Step 3 from the DS CLI.

Note: You can perform all steps from the DS CLI. The details are described in "Processing Remote FlashCopy (inband) transactions."

The following example illustrates the required steps for creating a remote FlashCopy operation.

- 1. **Create paths between the source LSS and the target LSS.** For example, IBM.2107-1300861 and IBM.2107-1300871 You need to know which volumes are available for use before you can issue the request to establish the path.
- 2. Create Metro Mirror volume pairs from the source LSS to the target LSS. For example, volume 2200 (IBM.2107-1300861/2200) from LSS22 and volume 2A00 (IBM.2107-1300871/2A00) from LSS22.
- 3. Enable a Remote FlashCopy operation at the target site using volume B as the source volume and volume C as the target volume. Assume that the target site network connection is lost. You can create the FlashCopy relationship from volume B to volume C (both volumes at the target site). However, you cannot use the DS8000 Storage Management GUI for this step because connections to the target site are lost.

Assume that you performed Step 1 and Step 2 from the DS8000 Storage Management GUI (connections to both storage units at the source and target sites were available at that time) and the Metro Mirror relationship between the volume pair still exists. To create the Remote FlashCopy operation, you must perform Step 3 from the DS CLI using the following command as an example. (You must be logged into the DS CLI in interactive command mode.)

Note: Use LSS 22 on the local site as a conduit LSS for the new remote Flash Copy relationship on the remote storage unit that will use volume 2A00 as the source. The target can be any other volume on the remote storage unit (in this example 2A01)

dscli> mkremoteflash -dev IBM.2107-1300871 -conduit IBM.2107-1300861/22 2A00:2A01

where:

-dev Specifies the storage image ID, which includes manufacturer, machine type, and serial number.

-conduit LSS_ID

(Required) Identifies the source LSS of an existing remote mirror and copy relationship that is used as a conduit for communicating with the target storage image. The source volume IDs that are specified in *source_volume_ID:target_volume_ID* must be the target volumes in a remote mirror and copy relationship in which one of the conduit LSS volumes acts as a source volume. You can specify a fully qualified LSS ID, which includes the storage image ID.

source_volume_ID:target_volume_ID

(Required) Establishes a remote FlashCopy relationship for the source and target volume pairs with the specified IDs. You can specify fully qualified volume IDs, which includes storage image IDs, or a shortened version without storage image IDs if the **-dev** parameter is specified. Separate the IDs of the FlashCopy relationships with spaces.

This report is displayed if your command input is correct.

FlashCopy Pair ID 2A00:2A01 successfully initiated. Use the lsremoteflash command to determine copy completion.

Verify that the transaction has processed successfully by issuing the following command:

dscli> lsremoteflash -dev IBM.2107-1300871 -conduit IBM.2107-1300861/22 2A00:2A01

Resynchronizing FlashCopy relationships

Complete this task to resynchronize (apply incremental changes on the source volume to) a FlashCopy target volume. After the initial FlashCopy operation, only data that has changed on the source volume since the last resynchronization operation was completed is copied to the target volume.

The change recording option and the persistent option must have been enabled on the FlashCopy volume pair. When a pair is established with the **-record** and **-persist** parameters, the pair initially synchronizes and then a record of all host write operations to the source is maintained in the source volumes.

You can resynchronize a FlashCopy target volume to create a point-in-time copy of your data without waiting to copy an entire volume for each point-in-time copy. Instead, only tracks that have changed on the source volume since the last resynchronization operation was completeed are copied to the target volume. The specified parameters in this command replace the parameters in the existing relationship. To keep the initial **-record** and **-persist** parameters, specify the **-record** and **-persist** parameters with the **resyncflash** command.

You can complete this task using either the DS CLI or the DS8000 Storage Management GUI.

Complete the following step to resynchronize FlashCopy relationships with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information the command requires. The second format is an example command with declared values for the variables.

Issue the **resyncflash** command to resynchronize FlashCopy relationships. Enter the **resyncflash** command at the dscli command prompt with the following parameters and variables:

dscli> resyncflash -dev storage_image_ID sourcevolumeID:targetvolumeID

Example dscli> resyncflash -dev IBM.2107-75FA120 0100:0200

The resulting output

FlashCopy pair 0100:0200 successfully incremented.

Reversing a FlashCopy relationship

Complete this task to reverse the direction of a FlashCopy volume pair.

When the direction of a FlashCopy relationship is reversed, 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, suppose 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 data on volume B is copied to volume A.

The background copy process must complete before you can reverse the direction of the FlashCopy relationship.

Exception: You cannot reverse the direction of the FlashCopy relationship during recovery from the failure of FlashCopy consistency group formation in a Global Mirror configuration due to a failure at the Global Mirror primary site. In this case, after you ensure the consistency of the FlashCopy consistency group target volumes, you can use the **fast** option of the **reverseflash** command before the background copy process completes to reverse the direction of the FlashCopy volume pair.

You can complete this task using either the DS CLI or the DS8000 Storage Management GUI.

Complete the following step to reverse the direction of FlashCopy relationships with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information the command requires. The second format is an example command with declared values for the variables.

Issue the **reverseflash** command to reverse the direction of FlashCopy relationships. Enter the **reverseflash** command at the dscli command prompt with the following parameters and variables:

dscli> reverseflash -dev storage_image_ID sourcevolumeID:targetvolumeID

Example

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

The resulting output

FlashCopy pair 0100:0200 successfully reverse restored.

Applying the FlashCopy revertible option to existing FlashCopy relationships

Complete this task to prepare for disaster recovery of a FlashCopy consistency group in a Global Mirror configuration. Issue the **setflashrevertible** command to a FlashCopy relationship with the persistent, change recording, target write inhibit, and no copy options enabled, and the revertible option disabled. It is not valid to issue the **setflashrevertible** command to a FlashCopy relationship that is already revertible.

The **setflashrevertible** command modifies a FlashCopy volume pair that is part of a Global Mirror relationship to *revertible*. If a failure occurs on the primary site during a Global Mirror create FlashCopy consistency group process and if that failure results in an inconsistency of the FlashCopy consistency group target volumes, you might be able to correct the inconsistency either by discarding changes or committing changes to the target volumes. The revertible option allows data to be committed to the target to form a new consistency group or to be reverted to the last consistency group. The **setflashrevertible** command must be run before the FlashCopy pair can be committed or reverted. You must have previously created a FlashCopy relationship with the persistent, change recording, target write inhibit, and no copy options enabled. The FlashCopy Revertible option must be disabled before beginning this task. It is not valid to complete the FlashCopy Revertible task on a FlashCopy relationship that is already revertible.

The FlashCopy Revertible task restarts an existing FlashCopy volume pair with the revertible option enabled for disaster recovery. The FlashCopy Revertible option remains in effect until the commit changes or discard changes task is completed. Both the commit changes and discard changes tasks disable the FlashCopy revertible option.

You can complete the FlashCopy Revertible task using either the DS CLI or the DS8000 Storage Management GUI.

Complete the following steps to apply the revertible option to existing FlashCopy relationships with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format is an example command with declared values for the variables.

Note:

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

Issue the **setflashrevertible** command to apply the revertible option to existing FlashCopy relationships. Enter the **setflashrevertible** command at the dscli command prompt with the following

parameters and variables: dscli> setflashrevertible -dev storage_image_ID sourcevolumeID:targetvolumeID

Note: Specify the storage unit for the **-dev** *storage_image_ID* parameter. This parameter is required if you do not specify a fully qualified ID for the source and target volumes and you do not specify a value for the *devid* variable in your profile file.

Example

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

The resulting output

FlashCopy volume pair 0100:0200 successfully made revertible.

Starting a background copy of a FlashCopy relationship

Complete this task to create a FlashCopy volume pair that allows data to be copied from the source volume to the target volume.

When you issue a FlashCopy command without using the **-nocp** parameter, the FlashCopy relationship is established but it is put in a queue for background copying. The exact time that the background copying starts for the specific relationship depends on the number of FlashCopy relationships that have already begun, or are waiting to begin, background copying. When the background copy starts, the state of that FlashCopy volume pair is displayed as "background copy running".

A background copy causes all data on the source volume to be physically copied to the target volume. After a FlashCopy pair is established, an automatic withdrawal of the FlashCopy relationship occurs when all source tracks have been physically copied to the target volume (unless the FlashCopy relationship was designated as persistent by using the **-persist** parameter when it was established).

Note: The amount of time that the actual physical copy can take depends on the amount of data that is copied and other activities that are occurring on the storage unit. You can monitor when the copy completes by issuing the **lsflash** command to check the status information for each FlashCopy relationship.

Complete the following steps to create FlashCopy relationships with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format is an example command with declared values for the variables.

1. Issue the **mkflash** command without the **-nocp** parameter to create FlashCopy relationships that allow data to be copied from the source volume to the target volume. Enter the **mkflash** command at the dscli command prompt with the following parameters and variables:

dscli> mkflash -dev storage_image_ID sourcevolumeID:targetvolumeID

Example

dscli> mkflash -dev IBM.2107-75FA150 0001:0004

Note: Specify the storage unit for the **-dev** *storage_image_ID* parameter. This parameter is required if you do not specify a fully qualified ID for the source and target volumes and you do not specify a value for the *devid* variable in your profile file.

A confirmation message is issued for each successful FlashCopy pair that is created.

2. Issue the **lsflash** command to check the status information for each FlashCopy relationship. A detailed report (when you use the **-1** parameter) is displayed for each FlashCopy relationship. Enter the **lsflash** command at the dscli command prompt with the following parameters and variables: dscli> lsflash -dev *storage_image_ID* -l *SourceVolumeID:TargetVolumeID*.

Example

dscli> lsflash -dev IBM.2107-75FA150 -1 0100:0100

Use the -1 parameter to provide a more detailed report about the FlashCopy relationships.

Preventing write operations on FlashCopy target volumes

Complete this task to prevent (inhibit) host write operations on FlashCopy target volumes. By inhibiting writes on the target volume, you ensure that the target is an uncorrupted incremental backup.

Use the **mkflash** command with the **-tgtinhibit** parameter to prevent host write operations on the target volume. When you use the **-tgtinhibit** parameter, the change recording feature is not active on the target volume. Write operations are not allowed on the target volume; therefore, the change recording bitmap for the target volume is not modified.

Note: By default, when you issue a **mkflash** command with the **-record** and **-persist** parameters, the FlashCopy relationship is established to act as an incremental FlashCopy. In addition, by default, when you issue the **setflashrevertible** command to a FlashCopy volume pair, the source volume of the volume pair is write inhibited. This allows the FlashCopy relationship to revert (change back) to a previous consistent state, if needed.

You can complete this task using either the DS CLI or the DS8000 Storage Management GUI.

Complete the following steps to prevent (inhibit) host write operations on FlashCopy target volumes with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information the command requires. The second format is an example command with declared values for the variables.

Issue the **mkflash** command with the **-tgtinhibit** parameter to prevent host write operations on the target volume of the FlashCopy relationships that you create. Enter the **mkflash** command at the dscli command prompt with the following parameters and variables: dscli> mkflash -dev *storage_image_ID* -tgtinhibit *sourcevolumeID:targetvolumeID*

Example

dscli> mkflash -dev IBM.2107-75FA150 -tgtinhibit 0001:0004

A confirmation message is issued for each successful FlashCopy pair that is created.

Creating a FlashCopy target volume on an existing Metro Mirror source volume

Complete this task to create a FlashCopy target volume on an existing Metro Mirror source volume.

Use the **mkflash** command with the **tgtpprc** parameter to create a FlashCopy target volume on an existing Metro Mirror source volume. The FlashCopy takes a point-in-time copy of a source volume, and then Metro Mirror makes a copy of the FlashCopy target volume at a remote site.

You can complete this task using either the DS CLI or the DS8000 Storage Management GUI.

Complete the following steps to create a FlashCopy target volume on an existing Metro Mirror source volume with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format is an example command with declared values for the variables.

• Issue the **mkflash** command with the **-tgtpprc** parameter to create a FlashCopy target volume on an existing Metro Mirror source volume. Enter the **mkflash** command at the dscli command prompt with the following parameters and variables:

dscli> mkflash -dev storage_image_ID -tgtpprc sourcevolumeID:targetvolumeID

Example

dscli> mkflash -dev IBM.2107-75FA150 -tgtpprc 0001:0004

A confirmation message is issued for each successful FlashCopy pair that is created.

Note: The **-tgtpprc** parameter can also be used with the **resyncflash** command. When you issue a **resyncflash** command to a FlashCopy relationship, only the new write operations to the source since the last resynchronization are copied to the target. This minimizes the data that is copied to the remote site when you also use the **-tgtpprc** parameter. The specified parameters in the **resyncflash** command replace the parameters in the existing relationship. To keep the initial **-record**, **-persist**, and **-tgtpprc** parameters must be specified in the **resyncflash** command.

• Issue the **resyncflash** command with the **-tgtpprc** parameter to resynchronize FlashCopy relationships and create a FlashCopy target volume on an existing Metro Mirror source volume. Enter the **resyncflash** command at the dscli command prompt with the following parameters and variables:

dscli> resyncflash -dev storage_image_ID -record -persist -tgtpprc sourcevolumeID:targetvolumeID

Example

dscli> resyncflash -dev IBM.2107-75FA120 -record -persist -tgtpprc 0100:0200

The resulting output

FlashCopy pair 0100:0200 successfully incremented.

Discarding changes to FlashCopy target volumes

Complete this task to discard changes to FlashCopy target volumes to form a consistency group on the target volumes as part of a disaster recovery process.

You cannot discard changes to FlashCopy target volumes unless you have modified the FlashCopy relationship using the **setflashrevertible** command, which changes the Revertible value to Enabled. You can use the **revertflash** command only when your analysis of the FlashCopy relationships reveals one of the following conditions:

- The FlashCopy relationships are revertible and all the sequence numbers are equal.
- There is a group of FlashCopy pairs that are all revertible and another group of FlashCopy pairs that are all nonrevertible. In addition, all the FlashCopy sequence numbers are not equal. However, the following conditions exist:
 - The FlashCopy sequence number for all revertible pairs is equal.
 - The FlashCopy sequence number for all nonrevertible pairs is equal.

If a FlashCopy consistency group formation operation does not complete, you must determine whether to discard changes (revert to a previous consistent state) or commit the operation to the current state. As part of a disaster recovery process, determine the state of the consistency groups in the affected sessions. The Discard Changes task specifies that the previous consistency group that was created by the Global Mirror session becomes the current state, and the Commit Changes task is no longer possible.

The Discard Changes task removes the FlashCopy relationship changes and resets them to the last consistency group state. The revertible state is set to No.

You can complete this task using either the DS CLI or the DS8000 Storage Management GUI.

Complete the following step to correct the applicable FlashCopy relationships with DS CLI commands. The example command in this task is shown in two formats. The first format shows the type of information that the command requires. The second format is an example command with declared values for the variables.

Issue the **revertflash** command to correct the FlashCopy relationships and reset them to the last consistency group state. Enter the **revertflash** command at the dscli command prompt with the following parameters and variables:

dscli> revertflash -dev storage_image_ID SourceVolumeID Example dscli> revertflash -dev IBM.2107-75FA150 0100

Notes:

- 1. Remember that *storage_image_ID* is the value for the remote server that has been designated the primary server until the original primary server is again available for use.
- 2. Global Mirror operations have completed the establish FlashCopy revertible processing but might have failed to form a consistency group before the disaster occurred. If your analysis, through use of the lsflash command, has determined that a revertflash command is needed, there is no need to issue a new mkflash command.

A confirmation message like the following one is generated for each FlashCopy relationship that has been successfully reset.

FlashCopy pair 0100:0200 successfully reverted to the previous consistency.

Committing data to FlashCopy target volumes

Complete this task to commit data to FlashCopy target volumes to form a consistency group on the target volumes as part of a disaster recovery process.

You can commit changes to FlashCopy target volumes only if you have modified the FlashCopy relationship using the **setflashrevertible** command, which changes the Revertible value to Enabled. You can use the **commitflash** command to commit data only when your analysis of the FlashCopy relationships reveals one of the following conditions:

- All the FlashCopy sequence numbers are equal and at least one of the FlashCopy relationships is nonrevertible.
- The FlashCopy relationships appear as follows:
 - Some of the FlashCopy relationships completed processing so that a consistent group was created. These FlashCopy relationships are no longer revertible.
 - Some of the FlashCopy relationships have not completed creating a consistency group. These FlashCopy relationships are still in a revertible state.
 - All the FlashCopy relationships have the same FlashCopy sequence number. This indicates that all the FlashCopy relationships are involved in the same consistency group.

If a FlashCopy consistency group formation operation does not complete, you must verify the consistency group at the remote site and determine whether the changes need to be "rolled forward" (committed) or "rolled backward" (discarded). The commit task specifies that the last consistency group that has been created by the Global Mirror session is committed to the current state, and reverting to the previous consistency group state is no longer possible.

You can complete this task using either the DS CLI or the DS8000 Storage Management GUI.

Complete the following step to correct the applicable FlashCopy relationships with DS CLI commands. The example command in this task is shown in two formats. The first format shows the type of information that the command requires. The second format is an example command with declared values for the variables.

Issue the **commitflash** command to correct the FlashCopy relationships and commit them to the consistency group that was being formed before the disaster occurred. Enter the **commitflash** command at the dscli command prompt with the following parameters and variables: dscli> commitflash -dev *storage_image_ID SourceVolumeID* **Example**

dscli> commitflash -dev IBM.2107-75FA150 0100

Note:

- Remember that *storage_image_ID* is the value for the remote server that has been designated the primary server until the original primary server is again available for use.
- Global Mirror operations have completed the establish FlashCopy revertible processing and might have failed to form a consistency group before the disaster occurred. If your analysis, through use of the **lsflash** command, has determined that a **commitflash** command is needed, there is no need to issue a new **mkflash** command.

A confirmation message like the following one is generated for each FlashCopy relationship that has been successfully reset.

FlashCopy pair 0100:0200 successfully committed.

Metro Mirror functions

Metro Mirror is a function for application data recovery, but also for failover to remote sites for disaster recovery, remote migration of data, and off-site backups. Refer to this topic for information that helps you use Metro Mirror functions when using the DS CLI commands.

Note: If you use Metro Mirror functions, you must disable the write acceleration feature for all switch manufacturers and models. The Global Mirror commands might fail if the write acceleration feature is enabled.

Displaying the status of established paths

Complete this task to display a list of established remote mirror and copy paths that are established between LSSs.

Before you begin with this task, ensure that the following guidelines are met:

- Fibre Channel I/O ports are configured.
- Fibre Channel paths have been established between source and target LSSs.

Use this task after you have issued the **mkpprcpath** command to determine the status of the paths that have been established between the specified source and target LSSs.

Complete the following step to display the status of established remote mirror and copy paths with DS CLI commands:

Issue the **lspprcpath** command to display the list of established paths. Enter the **lspprcpath** command at the dscli command prompt with the following parameters and variables: dscli> lspprcpath -dev *storage_image_ID Source_LSS_ID*. **Example**

dscli> lspprcpath -dev IBM.2107-75FA120 10

Note: You can specify multiple LSS IDs, but they must be separated with spaces.

Displaying the WWNN of a storage unit

Complete this task to display a list of worldwide node names (WWNNs) of the storage unit in a storage complex.

Before you begin, ensure that you have met the following conditions:

- The remote mirror and copy license key is installed and enabled to allow operations to be run.
- The Fibre Channel I/O ports are configured.

To participate in a Fibre Channel environment, each storage unit is assigned a unique 16-hexadecimal ID called a WWNN that identifies the storage unit. You must use the WWNN of the storage unit as part of the **lsavailpprcport** and **mkpprcpath** commands.

Complete the following steps to display the WWNN of the storage unit in a storage complex. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

- 1. Issue the **1ssi** command to display the list of WWNNs. Enter the **1ssi** command at the dscli command prompt as follows:
 - dscli> lssi -l
- **2**. Review the output that displays the WWNN of the storage unit. This information is required when you establish a path.

Creating remote mirror and copy paths

Complete this task when you create paths because they are required when creating source and target remote mirror and copy volume pair relationships.

Before you begin, ensure that you have met the following conditions:

- The remote mirror and copy license key is installed and enabled to allow operations to be run. If you are using a Model 2105 ESS as part of the configuration, ensure that you have the PPRC Version 2 license enabled.
- The I/O ports are configured for paths between source and target LSSs.
- The I/O ports to be used for paths are available and identified.
- The worldwide node name (WWNN) of the storage image is identified because it is a required parameter for this task.

Create paths so that the logical subsystems (LSSs) are associated with each other. Data is to be transferred through these paths so sufficient bandwidth for these operations is essential. In addition, you want to ensure that the ports used for remote mirror and copy operations are not the same ones to be used for host I/O activity.

Complete the following steps to create remote mirror and copy paths. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **mkpprcpath** command to create the Fibre Channel paths for the remote mirror and copy source and target volume pairs. Enter the **mkpprcpath** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprcpath -dev storage_image_ID -remotedev storage_image_ID -remotewwnn wwnn -srclss source_LSS_ID -tgtlss target_LSS_ID source_port_ID:target_port_ID

Notes:

- a. The -remotedev parameter specifies the ID of the secondary storage unit.
- b. The **-remotewwnn** parameter must specify the WWNN of the secondary storage unit. If you specify the WWNN of the primary storage unit, the command fails.
- c. You can specify the **-dev** and **-remotedev** parameters or specify fully qualified *srclss* and **-tgtlss** parameters, but not both.
- d. The shortened version of the -srclss and -tgtlss parameters are shown (value = 00) because the example uses the fully qualified -dev storage_image_ID parameter. If the fully qualified -dev parameter was not used, you must specify the fully qualified -srclss source_LSS_ID and the -tgtlss target_LSS_ID values. For example: -srclss IBM.2107-75FA120/00 (for DS8000), -srclss IBM.1750-75FA120/00 (for DS6000), -tgtlss IBM.2107-75FA120/01 (for DS8000), and -tgtlss IBM.1750-75FA120/01 (for DS6000).

e. The shortened version of the *source_port_ID:target_port_ID* parameter is shown (value = I1A10:I2A20), because the example uses the fully qualified -dev *storage_image_ID* and -remotedev *storage_image_ID* parameters. If the fully qualified -dev and -remotedev parameters are not used, you must use the fully qualified source_port_ID:target_port_ID value. For example: IBM.2107-75FA120/I1A10:IBM.2107-75FA150/I2A20.

The fully qualified *source_port_ID:target_port_ID* parameter is positional on your command line. It must be placed after the **-tgtlss** parameter. For example:

dscli> mkpprcpath -srclss 00 -tgtlss 00 IBM.2107-75FA120/I1A10:IBM.2107-75FA150/I2A20

Example

```
dscli> mkpprcpath -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150
-srclss 01 -tgtlss 00
-remotewwnn 12341234000A000F I1A10:I2A20
```

2. Issue the lspprcpath command to review the list of established remote mirror and copy paths.

Correcting a path-related configuration error

Complete this task to correct a path-related configuration error with DS CLI commands.

There might be occasions when you are using the **mkpprcpath** command to establish a path between the specified source and target LSSs and the transaction fails. You might not be aware of the failure until you run the **lspprcpath** command to check the status of the paths that have been established between the specified source and target LSSs.

The **lspprcpath** command displays a report that includes a *state* category. The state category reports on the current remote mirror and copy path state. One of the state codes is **configuration error**. A state code of configuration error is an indication that you have specified an incorrect value for the remote WWNN or the target lss ID.

Complete the following steps to correct the configuration error with DS CLI commands.

1. Check the original input values you provided for the -remotewwnn and -tgtlss parameters.

The following criteria applies to these parameters:

-remotewwnn

You must use the worldwide node name that is associated with the secondary storage unit. If you use the WWNN (worldwide node name) that is associated with the primary storage unit, the **mkpprcpath** command fails. Issue the **lssi** or **showsi** command to obtain the remote WWNN number of the secondary storage unit.

- **tgtlss** You must use the logical subsystem ID that is associated with the secondary storage unit as the target. You can verify that you have used the correct value by looking at the report that is generated by the **lspprcpath** command.
- 2. Obtain the correct values for the remote WWNN or target LSS ID and reissue the **mkpprcpath** command followed by issuing the **lspprcpath** command to verify that your transaction has processed correctly.

Removing paths

Complete this task to remove paths between the LSSs on the source storage unit and the target LSSs on the target storage units.

Before you delete paths, review the paths that are currently established.

If you delete all paths, you lose the communication between your remote mirror and copy volume pairs. All paths between the source LSS and target LSS are removed.

Complete the following steps to remove the paths between the source and target LSSs with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **lspprcpath** command to display a list of existing remote mirror and copy path definitions. Enter the **lspprcpath** command at the dscli command prompt with the following parameters and variables:

dscli> lspprcpath -dev storage_image_ID source_LSS_ID

Example

dscli> lspprcpath -dev IBM.2107-75FA120 01

The report that displays from this command provides the worldwide node name (WWNN) that is used with the **rmpprcpath** command.

2. Issue the **rmpprcpath** command to remove the paths between all source and target pairs. Enter the **rmpprcpath** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprcpath -dev storage_image_ID -remotedev storage_image_ID -remotewwnn wwnn source_LSS_ID:target_LSS_ID

Example

```
dscli> rmpprcpath -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -remotewwnn 12341234000A000F 01:01
```

Note:

- The **-remotedev** parameter specifies the ID of the secondary storage unit.
- The **-remotewwnn** parameter must specify the WWNN of the secondary storage unit. If you specify the WWNN of the primary storage unit, the command fails.
- If you do *not* specify the fully qualified **-dev** and **-remotedev** parameters, you must use the fully qualified *source_LSS_ID:target_LSS_ID* value. For example: IBM.2107-75FA120/01:IBM.2107-75FA150/01.

The fully qualified *source_LSS_ID:target_LSS_ID* value must be the last parameter in your command.

A confirmation message is displayed for each path that is being removed.

3. Enter **Y** to confirm that you want to remove the specified remote mirror and copy path. A message like the following example appears for each remote mirror and copy path that is being removed when you process the **rmpprcpath** command.

Are you sure you want to delete PPRC path (whatever was designated)? [y/n]: Y PPRC path (designated in the command) successfully deleted.

- PPRC path (designated in the command) successfully defeted.
- 4. Repeat Step 2 for all the remote mirror and copy paths that you want removed from the same source LSS to a different target LSS.

Creating a Metro Mirror relationship

Complete this task to create a Metro Mirror relationship between a source volume and target volume.

Before you begin, ensure that you have met the following conditions:

- The remote mirror and copy license key is installed and enabled to allow operations to run. If you are using a Model 2105 ESS as part of the configuration, ensure that you have enabled PPRC Version 2 license.
- The I/O ports are configured for paths between source and target LSSs.
- The Fibre Channel paths are created between all Metro Mirror source and target LSSs. The paths are required for communication between the volume pairs and to copy data from the source volumes to the target volumes. Otherwise, this task fails.

Metro Mirror is a function of a storage server that constantly updates a target copy of a volume to match changes made to a source volume. The source and target volumes can be on the same storage unit or on separate storage units. Metro Mirror creates the remote mirror and copy relationship in a synchronous manner.

Metro Mirror functions run on the DS8000 model or DS6000 model and are supported on many operating systems. For example, if you set up and configure your machine to use i5/OS, you can use Metro Mirror to create a copy of a System i disk pool on a separate machine, typically in a remote location.

Complete the following steps to create Metro Mirror relationships between the source volumes and target volumes.

- 1. Issue the **lsfbvol** command (for fixed blocked (FB) volumes) or the **lsckdvol** command (for count key data (CKD) volumes) to display which volumes are available for Metro Mirror relationships on the source and target LSSs. A report is displayed that shows the availability of the volumes.
- 2. Issue the **mkpprc** command to create a Metro Mirror relationship between a source volume and a target volume.

The **mkpprc** command must contain the following parameters and variables: dscli> mkpprc -dev *storage_image_ID* -remotedev *storage_image_ID* -type mmir *SourceVolumeID:TargetVolumeID*.

Note:

- The **-remotedev** parameter specifies the ID of the secondary storage unit.
- The **-type** *mmir* parameter specifies that you want to establish one or more Metro Mirror volume relationships. Metro Mirror creates the remote mirror and copy relationship in a synchronous manner.
- The shortened version of the *SourceVolumeID:TargetVolumeID* parameter is shown (value = 0100:0100) because the example uses the fully qualified -dev *storage_image_ID* and -remotedev *storage_image_ID* parameters. If the fully qualified -dev and -remotedev parameters were not used, you must use the fully qualified SourceVolumeID:TargetVolumeID value. For example IBM.2107-75FA120/0100:IBM.2107-75FA150/0100.

Example

```
dscli> mkpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -type mmir 0100:0100
```

A confirmation message is issued for each successful Metro Mirror volume association that is created.

3. Issue the **lspprc** command to view the status information of each Metro Mirror relationship in the list. Enter the **lspprc** command at the dscli command prompt with parameters and variables as follows: dscli> lspprc -dev *storage_image_ID* -remotedev *storage_image_ID* -s *SourceVolumeID:TargetVolumeID*. dscli> lspprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -s 0100:0100 0101:0101

Creating a Metro Mirror consistency group

Complete this task to create a Metro Mirror consistency group.

Ensure that the Remote Mirror and Copy license key is installed and enabled to allow the operations to run.

To restart applications at a remote site successfully, data at the remote site must be consistent. The Metro Mirror consistency group function keeps data consistency at the remote site by using consistency groups. A consistency group is defined a group of volumes that can temporarily queue (at the host level) subsequent write operations to all consistency group volumes on a single LSS pairing when an error occurs to one of the volumes in the group (source or target). A consistency group is also defined when a total link failure is detected between the source and target LSS volume pair.

The consistency group function of Metro Mirror consists of two parts. One is the consistency group option and the other is the freeze and unfreeze operation. The freeze operation causes the storage unit to suspend I/O activity through the **freezepprc** command. The unfreeze operation allows I/O activities to resume when you enter the **unfreezepprc** command.

Specify a consistency group option in either of the following two instances:

- When you define Metro Mirror paths between pairs of LSSs.
- When you change the default consistency group setting on each LSS, the consistency group option is disabled by default with the **chlss** command.

A group of volumes in a consistency group can consist of a combination of count-key-data volumes and fixed block volumes. A group of volumes in a consistency group can also consist of source volumes. These volumes can be associated with a DS6000 model or a DS8000 model and target volumes that are associated with an ESS 2105 Model 800 or 750.

Complete the following step to define a path that enabled the consistency group option for the volume pairs that are associated with the LSS volume pair. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with values declared for the variables.

1. Enter the **mkpprcpath** command to create a consistency group for the remote mirror and copy volume pairs. For example:

```
dscli> mkpprcpath -dev storage_image_ID
    -remotedev storage_image_ID -srclss source_LSS_ID -tgtlss target_LSS_ID
    -remotewwnn wwnn -consistgrp source_port_ID:target_port_ID
```

2. View the current consistency group state status of the consistency group by entering the **showlss** command. You can also use the **chlss** command to change the default consistency group timeout value. For example:

dscli> showlss -dev storage_image_ID LSS_ID

Resuming a Metro Mirror relationship

Complete this task to resume a Metro Mirror volume pair that was suspended (paused).

Before you begin, ensure that the following conditions are completed:

- The remote mirror and copy license key is installed and enabled to allow operations to be run. If you are using a Model 2105 ESS as part of the configuration, ensure that the PPRC Version 2 license is enabled.
- The Fibre Channel paths are created between all Metro Mirror source and target LSSs.

When you suspend (pause) volume pairs, Metro Mirror processing stops transferring data to the target volumes. Any I/O operations to the source volume are tracked during this time.

Use this task to resume a suspended (paused) Metro Mirror volume on the specified LSSs. When I/O is resumed, data is sent across to the target volumes.

Complete the following steps to resume Metro Mirror processing with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Enter the **resumepprc** command to continue Metro Mirror processing after it was suspended (paused).

dscli> resumepprc -dev storage_image_ID
-remotedev storage_image_ID -type [mmir, gcp] SourceVolumeID:TargetVolumeID

Notes:

1. The **-remotedev** parameter specifies the ID of the secondary storage unit.

- 2. Specify the **-type** parameter when you use the **resumepprc** command. Otherwise, the command fails.
- **3**. If you do not specify the fully qualified **-dev** and **-remotedev** parameters, you must use the fully qualified *SourceVolumeID:TargetVolumeID* value. For example: IBM.2107-75FA120/01:IBM.2107-75FA150/01 (DS8000) or IBM.1750-68FA120/01:IBM.1750-68FA150/01 (DS6000)

dscli> resumepprc -dev *IBM.2107-75FA120* -remotedev *IBM.2107-75FA150* -type mmir *0100:0100*

Pausing a Metro Mirror relationship

Complete this task to pause (suspend) a Metro Mirror relationship.

If you need to access target volumes or do maintenance on a remote storage unit, you can pause (or suspend) Metro Mirror volume pairs. This task pauses a Metro Mirror volume pair that you specify, and data is not copied to the target volume. The source storage unit keeps track of all changed data on the source volume, and after you resume the connection, only changes to the source volume are copied to the target volume.

Complete the following steps to pause Metro Mirror processing with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **pausepprc** command to pause Metro Mirror processing. Enter the **pausepprc** command at the dscli command prompt using the following parameters and variables:

dscli> pausepprc -dev storage_image_ID_-remotedev storage_image_ID SourceVolumeID:TargetVolumeID

Note:

- The **-remotedev** parameter specifies the ID of the secondary storage unit.
- If you do *not* specify the fully qualified **-dev** and **-remotedev** parameters, you must use the fully qualified **SourceVolumeID:TargetVolumeID** value. For example: IBM.2107-75FA120/01:IBM.2107-75FA150/01.

Example:

dscli> pausepprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 0100:0100

A confirmation message is displayed that indicates that processing for the specified volume pair has been paused.

After making your changes, you can resume processing by issuing the **resumepprc** command.

Creating a Global Copy relationship

Complete this task to create a Global Copy relationship between a source volume and target volume.

Before you begin, ensure that you have met the following conditions:

- The remote mirror and copy license key is installed and enabled to allow operations to run. If you are using a Model 2105 ESS as part of the configuration, ensure that you have PPRC Version 2 license enabled.
- The I/O ports are configured for paths between source and target LSSs.
- The Fibre Channel paths are created between all Metro Mirror source and target LSSs. The paths are needed for communication between the volume pairs and to copy data from the source volumes to the target volumes. Otherwise, this task fails.

You can create a Global Copy relationship between a source and target volume. Global Copy functions run on the DS8000 or DS6000 model and are supported on many operating systems. For example, if you set up and configure your machine to use i5/OS, you can use Global Copy to create a copy of a System i disk pool on a separate machine, typically in a remote location.

Complete the following steps to create Global Copy relationships between the source volumes and target volumes.

- 1. Issue the **lsfbvol** command (for fixed blocked (FB) volumes) or the **lsckdvol** command (for count key data (CKD) volumes) to display which volumes are available for Global Copy relationships on the source and target LSSs. A report is displayed that shows the availability of the volumes.
- 2. Issue the **mkpprc** command to create a Global Copy relationship between a source volume and a target volume. Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID -type gcp -mode full SourceVolumeID:TargetVolumeID

Notes:

- a. The -remotedev parameter specifies the ID of the secondary storage unit.
- b. The **-type** *gcp* parameter specifies that one or more Metro Mirror volume relationships be established. Global Copy creates the remote mirror and copy relationship in an asynchronous manner.
- c. The shortened version of the *SourceVolumeID:TargetVolumeID* parameter is shown (value = 0100:0100) because the example uses the fully qualified -dev *storage_image_ID* and -remotedev *storage_image_ID* parameters. If the fully qualified -dev and -remotedev parameters were not used, you must use the fully qualified SourceVolumeID:TargetVolumeID value. For example: IBM.2107-75FA120/0100:IBM.2107-75FA150/0100. This value must be placed at the end of your command line.

Example

dscli> mkpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -type gcp -mode full 0100:0100

A confirmation message is issued for each successful Global Copy volume association that is created.

3. Issue the **1spprc** command to view the status information of each Metro Mirror relationship in the list. Enter the **1spprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID -s SourceVolumeID:TargetVolumeID.
dscli> lspprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150
-s 0100:0100 0101:0101
```

Deleting a Metro Mirror relationship

Complete this task to delete the Metro Mirror relationship between a source and target volume.

You can use this task to delete the relationship between a Metro Mirror volume pair. The source and target volumes are removed from the configuration when this process runs.

The **rmpprc** command removes a remote mirror and copy (formerly PPRC) volume pair relationship.

Complete the following steps to delete the Metro Mirror relationship with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Enter the **lspprc** command to generate a report of Metro Mirror relationships. This can help you determine which Metro Mirror relationship that you want to delete.

dscli> lspprc -dev *storage_image_ID* -state -remotedev *storage_image_ID* SourceVolumeID:TargetVolumeID Example

dscli> lspprc -dev IBM.2107-75FA120 -l -remotedev IBM.2107-75FA150 0100: 01000101:0101

2. Enter the **rmpprc** command to delete the Metro Mirror relationship between the source and target volume.

dscli> rmpprc -dev storage_image_ID -remotedev SourceVolumeID:TargetVolumeID

Example

dscli> rmpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 0100:0100

Notes:

- a. If you delete a Metro Mirror volume pair with the source LSS and the process runs successfully, the source and the target volume go into the simplex state.
- b. If you delete a Metro Mirror volume pair with the target LSS and the process runs successfully, the source volume is in the suspended state, and the target volume is in the simplex state. This option is useful in a disaster situation when the source (local) site failed.

Modifying logical subsystem timeout values

This section describes the list of logical subsystem (LSS) timeout values that you can modify using DS CLI commands.

The following lists contains the LSS timeout values that you can modify using the **mkpprcpath** command:

- Concurrent copy timeout (zSeries only)
- Consistency group timeout
- Critical mode enable (zSeries only)
- z/OS Global Mirror timeout (DS8000 only)

Modifying the Concurrent Copy timeout value

Complete this task to modify a Concurrent Copy timeout value, which determines how long a volume in a Concurrent Copy session stays "long busy" (unavailable) before suspending the session. This topic applies to z Systems only.

Use the **chlcu** command to modify the Concurrent Timeout value of the logical control unit. The Concurrent timeout value determines how long a logical volume in a Concurrent Copy session in a specified LSS remains in a long-busy condition before the session is suspended.

Complete the following step to modify a Concurrent Copy timeout value (the DS CLI command refers to this as **-ccsess** *timeout* or Concurrent Copy session timeout). The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **chlcu** command to modify a Concurrent Copy timeout value. Enter the **chlcu** command at the dscli command prompt with the following parameters and variables: dscli> chlcu -dev *storage_image_ID* -ccsess *timeout LCU_ID* **Example**

dscli> chlcu -dev IBM.2107-75FA120 -ccsess 190 00-0F

A confirmation message is displayed for each LCU that has been modified.

Modifying the consistency group timeout value

Complete this task to modify the consistency group timeout value, which determines the amount of time that I/O is withheld from updating a source volume of a consistency group after an error occurs.

The consistency group timeout value (the **-extlongbusy** *timeout* parameter) is the time in seconds that a volume in a Metro Mirror consistency group stays unavailable after an error causes the suspension of a consistency group operation if a consistency group is not received before the timeout value. The consistency group timeout value enables automation software to detect that an error has occurred and to issue a command to freeze all other volumes of the consistency group. When an error is detected, a long-busy condition occurs.

Complete the following step to modify the consistency group timeout value. The example command in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **chlcu** command to modify the consistency group timeout value. Enter the **chlcu** command at the dscli command prompt with the following parameters and variables: dscli> chlcu -dev *storage_image_ID* -extlongbusy *timeout LCU_ID* **Example**

dscli> chlcu -dev IBM.2107-75FA120 -extlongbusy 3 00-0F

Modifying the z/OS Global Mirror timeout value

Complete this task to modify the time that a volume in an z/OS Global Mirror session remains in a long-busy condition (the volume is not available) before the session is suspended. This task can be done in the DS8000 only.

You can modify the z/OS Global Mirror timeout value, which determines how long any volume in the selected LSS in a z/OS Global Mirror session remains long busy before the session is suspended. The default value is 300 seconds. (The z/OS Global Mirror timeout value is also known as the Extended Remote Copy (XRC) timeout value that is supported on the ESS.)

The long-busy condition occurs because the system data mover cannot copy data when the volume (or z/OS Global Mirror session) is not able to accept additional data. The value of this timeout is associated with a z/OS Global Mirror session when the session is created.

Complete the following step to modify the z/OS Global Mirror timeout value. The example command in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with values declared for the variables.

Issue the **chlcu** command to modify the z/OS Global Mirror timeout value. Enter the **chlcu** command at the dscli command prompt with the following parameters and variables: dscli> chlcu -dev *storage_image_ID* -xrcsess *timeout LCU_ID* **Example**

dscli> chlcu -dev IBM.2107-75FA120 -xrcsess 175 00-0F

Modifying the critical mode setting

Complete this task to enable the critical mode setting to prevent write operations to source volumes if data cannot be copied to the target volume of the volume pair because of a permanent error. You must have administrator authority to complete this task.

The critical mode setting is used to determine the behavior of remote mirror and copy (PPRC) pairs or consistency groups after the source and target storage units can no longer communicate or when paths between a volume pair in the specified LSS are lost. This setting is associated with the volume pairs in the LSSs that you selected. This option is available for z/OS environments only.

When you enable the critical mode setting, the volume pair is suspended and further write operations to the source volume are not accepted if data cannot be sent to the target volume. The volume pair remains in a suspended state until you correct the problem and either issue a request to resynchronize the volume pair or delete it.

If you do not enable this setting and an error occurs to the target volume, the remote mirror and copy feature suspends the copy pair, which allows the subsequent write operations to be copied to the source volume of that volume pair. The storage unit records all tracks that have changed. When the problem is resolved, you can resynchronize the volume pair.

Complete the following step to modify the critical mode setting. The example command in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **chlcu** command to modify the critical mode setting. Enter the **chlcu** command at the dscli command prompt with the following parameters and variables: dscli> chlcu -dev *storage_image_ID* -critmode *enable LCU_ID* **Example**

dscli> chlcu -dev IBM.2107-75FA120 -critmode enable 00-0F

Note: 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.

This parameter cannot be used with Global Copy or remote copy and mirror cascading volumes. This parameter only applies to S/390[®] or z Systems volumes.

The following table presents an overview of how the critical volume mode works.

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 primary volume when the last path to the secondary volume has failed. Inhibits write operations to the primary volume.
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.

Defining a path that has the consistency option enabled

Complete this task to define a path that has enabled the consistency group option for the volume pairs that are associated with the LSS volume pair.

Ensure that the Remote Mirror and Copy license key is installed and enabled to allow the operations to run.

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. This is the only supported connectivity for machine types 2107 and 1750. Paths can be established between the following machine types: 2105:2105, 2107:2107, 2107:1750, 2107:2105, 1750:2105.

A consistency group is a group of volumes that provides the ability to temporarily queue (at the host level) subsequent write operations to all consistency group volumes on a single LSS pairing when an error occurs to one of the volumes in the group (source or target), or when a total link failure is detected between the source and target LSS volume pair.

This process describes how to define paths that have enabled the consistency group option. This means that when an error occurs on any volume pairs or on the links that are associated with these LSS pairs, an alert is issued and I/O to all duplex remote mirror and copy volumes on LSS pairs will be queued either until a consistency group creation operation is run or the consistency group timeout time expires. This allows external automation to use the consistency group created operation to create a dependent write consistent set of target volumes over any number of LSS and disk storage units.

Complete the following step to define a path that has enabled the consistency group option for the volume pairs that are associated with the LSS volume pair. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with values declared for the variables.

Issue the **mkpprcpath** command to create a consistency group for the remote mirror and copy volume pairs. Enter the **mkpprcpath** command with the **-consistgrp** parameter at the dscli command prompt with the following parameters and variables:

dscli> mkpprcpath -dev *storage_image_ID* -remotedev *storage_image_ID* -srclss *source_LSS_ID* -tgtlss *target_LSS_ID* -remotewwnn *wwnn* -consistgrp *source_port_ID:target_port_ID* **Example**

dscli> mkpprcpath -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -srclss 01 -tgtlss 01 -remotewwnn 12341234000A000F -consistgrp I0100:I0100

Monitoring Remote Mirror and Copy paths

Complete this task to display a list of existing remote mirror and copy path definitions using the DS CLI.

Complete the following step to display a list of existing remote mirror and copy path definitions. The example command in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **lspprcpath** command to generate a report of existing remote mirror and copy path definitions. Enter the **lspprcpath** command at the dscli command prompt with the following parameters and variables:

dscli> lspprcpath -dev *storage_image_ID* -fullid *Source_LSS_ID* **Example**

dscli> lspprcpath -dev IBM.2107-75FA120 -fullid 10

Running a recovery failback operation

Complete this task to run Global Copy failback processing for the A volumes. This process resynchronizes the volumes at Site A with volumes at Site B and restarts mirroring from site A (local site) to site B (remote site).

You must first create a remote mirror and copy volume pair. Before you run the failback operation, the volumes must be full duplex.

The **failbackpprc** command copies the required data from the source volume to the target volume to resume mirroring. The failover process converted the full-duplex target volumes at site A to suspended source volumes. The volumes at site A started the change recording process while in failover mode.

The failback processing that is described in this task can be issued 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.

Complete the following step to run a recovery failback operation. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **failbackpprc** command to run a recovery failback operation. Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev *storage_image_ID* -remotedev *storage_image_ID* SourceVolumeID:TargetVolumeID **Example**

```
dscli> failbackpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 0100-0103:0100-0103
```

2. Issue the **1spprc** command to check the status information of each Metro Mirror relationship in the list. Enter the **1spprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID -s
SourceVolumeID:TargetVolumeID
```

Example

```
dscli> lspprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150
-s 0100:0100 0200:0200 0300:0300
```

Running a recovery failover operation

Complete this task to run a recovery failover operation using DS CLI Metro Mirror. In a disaster recovery process, the failover procedure must be followed by a failback procedure after a path from the target site to the source site is created.

Create a remote mirror and copy volume pair. Volume sizes for operations that use failover and failback operations must be the same; otherwise, the failback operation fails.

A failover to the Global Copy secondary volume turns the secondary volumes into primary volumes and suspends these volumes immediately. When you run a Global Copy failover, the B volumes are the primary volumes and the A volumes are the secondary volumes. This action just changes 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, but leaves 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.

Complete the following step to run a failover recovery operation. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **failoverpprc** command to run a recovery failover operation. Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables: dscli> failoverpprc -dev *storage_image_ID* -remotedev *storage_image_ID SourceVolumeID:TargetVolumeID* **Example**

dscli> failoverpprc -dev IBM.2107-75FA150 -remotedev IBM.2107-75FA120 0100-0103:0100-0103

Viewing information about Metro Mirror relationships

Complete this task to view information about Metro Mirror relationships using the DS CLI.

The **lspprc** command displays a list of remote mirror and copy (formerly PPRC) volume relationships for a storage image (or storage unit for DS6000), and status information for each remote mirror and copy volume relationship in the list.

Complete the following step to view information about Metro Mirror relationships. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **lspprc** command to generate a report of Metro Mirror relationships. Enter the **lspprc** command at the dscli command prompt with the following parameters and variables: dscli> lspprc -dev *storage_image_ID* -state -remotedev *storage_image_ID* SourceVolumeID:TargetVolumeID

Example

dscli> lspprc -dev IBM.2107-75FA120 -1 -remotedev IBM.2107-75FA150 0100:0100 0101:0101

Converting Global Copy volume pairs to synchronous

Complete this task to convert Global Copy volume pairs to synchronous (Metro Mirror volume pairs).

Before you begin, ensure that the license for the remote mirror and copy feature is activated. Paths are required between the source and the target LSS storage units for the volume pairs.

There are two common situations where you would convert a Global Copy volume pair to a Metro Mirror volume pair:

- You have used the Global Copy function to complete the bulk transfer of data in the creation of many copy pairs, and you now want to convert some or all of those pairs to Metro Mirror mode.
- You have Global Copy pairs for which you want to make FlashCopy backups on the remote site. You convert the pairs temporarily to synchronous mode to obtain a point-in-time consistent copy.

If you created a Global Copy volume pair where the source volume was associated with a DS8000 or a DS6000 machine type and the target volume was associated with an ESS 2105 Model 800 or 750, you can convert that volume pair to synchronous.

Complete the following step to convert Global Copy volume pairs to synchronous (Metro Mirror volume pairs). The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **mkpprc** command to convert Global Copy volume pairs to synchronous (Metro Mirror volume pairs). Enter the **mkpprc** command with the **-type** *mmir* parameter at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID -type mmir
SourceVolumeID:TargetVolumeID
```

Example

dscli> mkpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 0100:0100 -type mmir 0101:0101 0102:0102 0103:0103

Determining which I/O ports are available for paths

Complete this task to determine which I/O ports are available for paths between the source and target LSSs.

Before you begin with this task, ensure that the following guidelines are met:

- The remote mirror and copy license key is installed and enabled to allow operations to be run.
- The Fibre Channel I/O ports are configured.
- The I/O ports that will be used for paths are available and identified.

• The WWNN of the storage unit is identified because it is a required parameter for this task.

Before you create paths, use this task to determine which ports are available for remote mirror and copy (formerly PPRC) I/O operations. These are the ports through which data will be transferred so it essential that bandwidth for these operations be sufficient. In addition, you want to ensure that the ports used for remote mirror and copy operations are not the same ones that will be used for host I/O activity.

You need to determine which source and target I/O ports are available for paths on the local and remote storage units. The output that is generated from this task displays ESCON (DS8000 only) or Fibre Channel protocol (FCP) I/O ports that are available to be used as remote mirror and copy paths. The Enterprise Storage Server (2105 machine type) supports ESCON ports (DS8000 only).

Note: When you establish FCP paths, the LSSs on the source and target storage units can be connected either through a point-to-point connection (no switch) or through a switched fabric. For Fibre Channel attachments, you can establish zones to help reduce the possibility of interactions between system adapters in switched configurations. For information, see the Fibre Channel switches publication that is available for your environment.

Complete the following steps to determine the available I/O ports with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with values declared for the variables.

1. Issue the **lsavailpprcport** command to display a list of available I/O ports that are available for paths. Enter the **lsavailpprcport** command at the dscli command prompt with the parameters and variables shown as follows:

dscli> lsavailpprcport -dev storage_image_ID -remotedev storage_image_ID -remotewwnn wwnn source_LSS_ID:target_LSS_ID

Notes:

- a. The -remotedev parameter specifies the ID of the secondary storage unit.
- b. The **-remotewwnn** parameter must specify the worldwide node name of the secondary storage unit. If you make a mistake and specify the worldwide node name of the primary storage unit, the command fails.
- c. The shortened version of the *source_LSS_ID:target_LSS_ID* parameter is shown (value = 01:01) because the example uses the fully qualified -dev *storage_image_ID* and -remotedev *storage_image_ID* parameters. If the fully qualified -dev and -remotedev parameters were not used, you must use the fully qualified source_LSS_ID:target_LSS_ID value. For DS8000, example: IBM.2107-75FA120/01:IBM.2107-75FA150/01

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
```

Example

dscli> lsavailpprcport -l -dev IBM.2107-75FA120
 -remotedev IBM.2107-75FA150 -remotewwnn 12341234000A000F 01:01

2. Analyze the output that is generated and select from the available I/O ports to create the path. The information that is displayed shows available I/O ports combinations between the source LSSs and the target LSSs and the output depends on the current selection of adapters.

Deleting a single volume Metro Mirror relationship

Complete this task to delete the single volume Metro Mirror relationship whether it exists on the source or the target volume.

There might be times when a communication problem occurs between your primary server and your secondary server or vice versa. If this problem happens during the processing of a **rmpprc** command, only part of the removal transaction is completed. The removal takes place even though the error message might indicate that the entire process failed.

To correct this situation and to remove the other part of the pair relationship, you must reenter the **rmpprc** command for each volume that was not removed. Use the following parameters:

- The -at src parameter, if the pair relationship was not removed from the source volumes.
- The **-at***tgt* parameter, if the pair relationship was not removed from the target volumes.
- The **-unconditional** parameter with the **-at***src* or **-at***tgt* parameter; otherwise, the transaction fails.

When the transaction completes, the affected volume is returned to a simplex state and becomes available for use in another pair relationship.

Complete the following steps to delete the Metro Mirror relationship with DS CLI commands. The example commands in this task are shown in two formats. The first format shows the type of information the command requires. The second format provides the command with declared values for the variables.

1. Enter the **lspprc** command to generate a report about Metro Mirror relationships. The command output can help you determine which Metro Mirror relationships must be deleted.

```
dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID -l
SourceVolumeID:TargetVolumeID
```

Example

dscli> lspprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -1 0100:0100 0101:0101

2. Enter the **rmpprc** command to delete the pair relationship that the volume still maintains.

Source volume

```
dscli> rmpprc -dev storage_image_ID -at src -unconditional
SourceVolumeID
```

Example

```
dscli> rmpprc -dev IBM.2107-75FA120 -at src -unconditional 0100
```

Target volume

```
dscli> rmpprc
-dev storage image ID -at tgt -unconditional storage image ID
```

Example

dscli> rmpprc -dev IBM.2107-75FA150 -at tgt -unconditional 0100

Notes:

- a. The **-dev** parameter must contain the value of the secondary server when you are removing the pair relationship from a target volume.
- b. The management console must be able to communicate with the secondary server for this command to succeed.

Copy Services functions across a 2105 and DS8000 or DS6000 model

Copy Services functions that are run using either the DS8000 Storage Management GUI or DS CLI can interact with either DS8000 models or the DS6000 model and the IBM 2105 Enterprise Storage Servers (ESS) Models 800 and ESS 750.

Most Copy Services functions that are available on the ESS 2105 are also functional on DS8000 models or the DS6000 model and in open systems and z Systems environments.

On the DS8000 models, a storage unit image consists of two logical partitions (LPARs), one LPAR on each processor complex. Each LPAR is allocated resources (processors, memory, I/O adapters and storage devices). The combination of the two LPARs creates an image of a single DS8000 machine which can be seen by the 2105 machine. You can use the mirroring solutions that are compatible between the 2105 and DS8000 to set up disaster recovery solutions with either machine being the primary (local) site, or the secondary (remote) site.

Before you begin, consider the following guidelines:

- To run Copy Services functions between machine types that involve a 2105, you must configure a Copy Services domain on the DS8000 Storage Management GUI or DS CLI.
- To connect to the 2105 Copy Services domain on the ESS, all interfaces that you use require an authenticated login procedure to access Copy Services functions across the storage complex. The authentication is completed by using a user name and password that was created with the ESS Specialist. Therefore, the existing user name and password that was created with the ESS Specialist for the 2105 Copy Services domain for which you will be working must match the user name and password on the management console that is connected to the DS8000 or DS6000 model. Otherwise, you must add them using either the DS8000 Storage Management GUI or DS CLI as part of the procedure for adding a 2105 Copy Services domain to the storage complex.
- To manage Copy Services across the 2105 and a DS8000 or DS6000 model, you must install licensed internal code version 2.4.2 or later on the 2105.
- The 2105 and DS6000 model does not support remote mirror and copy (formerly PPRC) operations with an ESCON link. If you want to configure a remote mirror and copy relationship between the DS8000 or DS6000 model and a 2105, you must use a FCP link.

Creating a Metro Mirror volume pair between a DS8000 model or a DS6000 model and a 2105

Complete this task to create a Metro Mirror volume pair using volumes from a DS8000 or a DS6000 model and a 2105.

Before you begin, ensure that you meet the following requirements:

- The license for the remote mirror and copy feature must be activated.
- To create a Metro Mirror volume pair between DS8000 or DS6000 models and a 2105, you must have added the 2105 Copy Services domain to your storage complex environment.
- Ensure that paths are set up between the source and the target LSSs for the Metro Mirror volume pairs. The paths between the 2105 and the DS8000 or the DS6000 model must be configured using Fibre Channel Protocol (FCP) ports.
- The storage type of the source and target volumes on the DS8000 or the DS6000 model and 2105 domain must have the same type. That is, if the source volumes are fixed block volumes, the target volumes must also be fixed block volumes.
- The size of the volumes in the source LSS must be less than or equal to those of the target LSS.
- Gather the following preliminary information:
 - Open the ESS Specialist on the ESS 800 to determine the its WWNN. The WWNN is listed in 20 point font on the opening page. The format is **5005076300C08641**.
 - Determine the number of available volumes on the ESS 800 with the ESS 800 GUI.
 - Document the LSS and volume mappings.
 - Ensure that the volume sizes are matched and are -type ESS on the DS6000 or the DS8000.

You can create Metro Mirror relationships using source and target volumes from the following machine types:

- A DS8000 and a DS8000 model
- A DS6000 and a DS6000 model

- A DS8000 and a DS6000 model
- A 2105 and a DS6000 model
- A 2105 and a DS8000 model
- A 2105 and a 2105

Note: If the source is a Copy Services 2105 domain, the Metro Mirror task is completed on the source domain. However, if you run a "Suspend at target" action, the suspension occurs at the target domain.

Complete the following steps to create a Metro Mirror pair between a DS8000 or a DS6000 model and a 2105. For this task, the source domain is a 2105 Model 800 or 750 and the target is a DS8000 or a DS6000 model. You can use this task if the target domain is a 2105 Model 800 or 750 and the source is a DS8000 or a DS6000 model by switching the device IDs in the volume pairs. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **lsavailpprcport** command to list all of the available ports to the remote system to create a connection. Ensure that you use ports that are not mapped to hosts for PPRC for increased performance, but sharing host ports with PPRC ports is supported. Specify the remote WWPN and device ID for the target cluster. Enter the **lsavailpprcport** command at the dscli command prompt with the following parameters and variables:

dscli> lsavailpprcport -dev storage_image_ID -remotedev storage_image_ID -remotewwnn wwnn Source_LSS_ID:Target_LSS_ID

Example

dscli> lsavailpprcport -dev IBM.2107-13AB7DA -remotedev IBM.2105-18602 -remotewwnn 5005076300C08642 10:10

2. Issue the **mkpprcpath** command to create a path between LSSs on the DS8000 or the DS6000 model to the ESS 800. You associate an LSS on the DS8000 or the DS6000 model to the ESS 800 and specify specific ports. You can list multiple ports. You should create redundant port paths from both controllers of the DS8000 or the DS6000 model to both clusters of the ESS 800. Enter the **mkpprcpath** command at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprcpath -dev storage_image_ID -remotedev storage_image_ID
-srclss source_LSS_ID -tgtlss target_LSS_ID -remotewwnn wwnn
-consistgrp source_port_ID:target_port_ID
```

Example

dscli> mkpprcpath -dev IBM.2107-75FA120 -remotedev IBM.2105-18602 -srclss 01 -tgtlss 01 -remotewwnn 12341234000A000F -consistgrp I0100:I0100

3. Issue the **1spprcpath** command to display the created paths and their status. "Success" indicates that the path is valid, "failure" indicates that the path did not create correctly, or that the relationship has become separated. Enter the **1spprcpath** command at the dscli command prompt with the following parameters and variables:

dscli> lspprcpath -dev storage_image_ID Source_LSS_ID Example

dscli> lspprcpath -dev IBM.2107-75FA120 10

4. Issue the **rmpprcpath** command to remove a path. You must specify both the source and target device IDs and the source and destination LSS. Enter the **rmpprcpath** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprcpath -dev storage_image_ID -remotedev storage_image_ID source_LSS_ID:target_LSS_ID

Example

dscli> rmpprcpath -dev IBM.2107-75FA120 -remotedev IBM.2105-18602 10:10

5. Issue the **mkpprc** command to create a relationship between a source and target volume. The volumes must be type ESS and be exactly the same size or the attempt to create the volume pair will fail. Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID -type mmir SourceVolumeID:TargetVolumeID

Example

dscli> mkpprc -dev IBM.2107-75FA120 -remotedev IBM.2105-18602 -type mmir 1001:1001

6. Issue the **lspprc** command to list the pairs that are in existence. Upon creation, the volumes will be in the *copy pending* state. When the initial copy is complete, the volumes will show as *full duplex* on the primary and *target full duplex* on the secondary. If something interrupts the connection, the primary volumes indicate *suspended*, but the target volumes still show *full duplex*. You can specify a range of volumes to list multiple pairs 1001 - 10ff. Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev storage_image_ID SourceVolumeID:TargetVolumeID

Example

dscli> lspprc -dev IBM.2107-75FA120 1001-10ff

Global Mirror functions

Global Mirror asynchronously copies data from a host to a remote site, and maintains a consistent copy of the data on a storage unit at the remote site. Refer to this topic for information that helps you use Global Mirror functions when using the DS CLI commands.

Note: If you use Global Mirror functions, you must disable the write acceleration feature for all switch manufacturers and models. The Global Mirror commands might fail if the write acceleration feature is enabled.

Adding volumes to a session (Global Mirror)

Complete this task to add volumes to a Global Mirror session.

You can add Global Copy primary volumes to a Global Mirror session at any time after the Global Mirror session has started without stopping the session. If you attempt to add a Metro Mirror volume or volumes which, for example, is converted from Global Copy to Metro Mirror, the formation of the consistency group fails.

Volumes can be added to a Global Mirror session but do not become active in the session until the Global Copy pair has completed its first pass and a consistent copy of the data has been formed at the remote site.

If you have many volumes that you want to add to a Global Mirror session, you might consider adding them to the session in stages. This lessens the impact on your processing.

Complete the following steps to add volumes to a Global Mirror session. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

 Issue the lspprc command to obtain a list of the Global Copy volumes that you can add to the Global Mirror session. A detailed report is displayed (if you use the -l parameter) that allows you to see the available volumes. Enter the lspprc command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID -l SourceVolumeID:TargetVolumeID.

Example

dscli> lspprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -1 0100:0100

 Issue the chsession command to add the available volumes to a Global Mirror session. Enter the chsession command at the dscli command prompt with the following parameters and variables: dscli> chsession -dev storage_image_ID -lss LSS_ID -action add -volume volume_ID session_ID

Example

dscli> chsession -dev IBM.2107-75FA120 -lss 10 -action add -volume 0100-010F 01

A confirmation message indicates that the session has been modified successfully.

3. Issue the **1ssession** command to query the status of all volumes being processed including the volumes that you added to the session. Enter the **1ssession** command at the dscli command prompt with the following parameters and variables:

dscli> lssession -dev storage_image_ID -l LSS_ID

dscli> lssession -dev IBM.2107-75FA120 -1 01

When you use the **-1** parameter, a detailed report displays a list of Global Mirror sessions for the specified logical subsystem (LSS) and information about the volumes of each session in the list.

Modifying the tuning parameters of a Global Mirror session

Complete this task to modify the tuning parameters of a Global Mirror session.

A global mirror session consists of tuning parameters and topology, both of which can be modified. However, they cannot be modified using the same method. The modification of the tuning parameters requires that you pause the Global Mirror session and change the parameters. Then you can resume the session with the new tuning parameter values. The modification of the Global Mirror topology requires that you stop the Global Mirror session, change the topology, and then restart Global Mirror processing.

A Global Mirror session includes the following tuning parameters:

- · Consistency group interval time
- Maximum coordination interval time
- Maximum consistency drain time

There are two specific occasions when you might want to modify the tuning parameters:

- After the initial setup of your session and your analysis of Global Mirror processing indicates that these parameters should be adjusted.
- After you have stopped (not just paused) Global Mirror processing. When you restart Global Mirror processing, the values for the tuning parameters revert to their DS CLI default values.

Complete the following steps to modify the tuning parameters of a Global Mirror session. The example commands that are displayed in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **pausegmir** command to pause Global Mirror processing on the specified logical subsystem and the specified session within the logical subsystem. Enter the **pausegmir** command at the dscli command prompt using the following parameters and variables:

dscli> pausegmir -dev storage_image_ID -lss LSS_ID -session session_ID Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID

Example

dscli> pausegmir -dev IBM.2107-75FA120 -lss 10 -session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00

A confirmation message indicates that the session has been paused after all buffered write operations to the target have been processed.

2. Issue the **showgmir** command to receive a report that provides the current properties or performance metrics for a Global mirror logical subsystem ID. Enter the **showgmir** command at the dscli command prompt with the following parameters and variables:

For a detailed properties report: dscli> showgmir -dev storage_image_ID LSS_ID

For a performance metrics report: dscli> showgmir -dev storage_image_ID -metrics LSS_ID

Example

These commands are entered as follows when you add values:

dscli> showgmir -dev IBM.2107-75FA120 10

dscli> showgmir -dev IBM.2107-75FA120 -metrics 10

- **3**. Analyze the report and determine which if any of the Global Mirror tuning parameters must be changed.
- 4. Issue the **resumegmir** command with the values for all three tuning parameters. Enter the **resumegmir** command at the dscli command prompt with the following parameters and variables:

dscli> resumegmir -dev storage_image_ID -lss ID -cginterval seconds -coordinate milliseconds -drain seconds -session_ID Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID

Example

The example command shows all three tuning parameters with new values. You must specify a value for all three tuning parameters even if only one value has changed. The values for the two unchanged tuning parameters would be the DS CLI default values.

dscli> resumegmir -dev IBM.2107-75FA120 -lss 10 -cginterval 10 -coordinate 60 -drain 35 -session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00

Notes:

- a. The **-cginterval** parameter specifies the consistency group interval time, in seconds. The value specifies how long to wait between the formation of consistency groups. If this value is set to zero, consistency groups are formed continuously. The DS CLI default value is 0.
- b. The **-coordinate** parameter 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 DS CLI default value is 50 milliseconds.
- **c.** The **-drain** parameter specifies the maximum amount of time that writes (in seconds) are inhibited to the remote site before the current consistency group must stop. The DS CLI default value is 30 seconds.

Modifying the topology of a Global Mirror session

Complete this task to modify the topology of a Global Mirror session that is not part of a script.

A Global Mirror session consists of tuning parameters and topology, both of which can be modified. However, they cannot be modified using the same method. The modification of the Global Mirror topology requires that you stop the Global Mirror session, change the topology, and then restart Global Mirror processing. The modification of the tuning parameters is managed differently.

Topology in this process refers to the list of subordinate storage servers. You establish remote mirror and copy paths between the master and subordinate LSSs. Just one LSS per subordinate server is sufficient. When you define the remote mirror and copy path, you identify the primary LSS on the master server. The secondary LSS in the remote mirror and copy path establishes command points to a corresponding subordinate server. These LSSs are part of the topology specification that defines the communication paths between the master and subordinate storage servers.

Note: When you restart Global Mirror processing, the tuning parameters revert to their DS CLI default values.

Complete the following steps to modify the topology of a Global Mirror session. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **showgmir** command to display a detailed properties or performance metrics report for a Global Mirror logical subsystem ID. Enter the **showgmir** command at the dscli command prompt with the following parameters and variables:

For a detailed properties report: dscli> showgmir -dev storage_image_ID -fullid LSS_ID

For a performance metrics report: dscli> showgmir -dev storage_image_ID -metrics LSS_ID

These commands are entered as follows when you add values:

dscli> showgmir -dev IBM.2107-75FA120 -fullid 10

dscli> showgmir -dev IBM.2107-75FA120 -metrics 10

- 2. Use the report as a guide to see what is currently being processed and to determine what topology values you want to change.
- **3**. Issue the **rmgmir** command to stop Global Mirror processing. Enter the **rmgmir** command at the dscli command prompt with the following parameters and variables:

dscli> rmgmir -dev storage_image_ID -lss ID -session session_ID Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID

Example

dscli> rmgmir -dev IBM.2107-75FA120 -lss 10 -session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00

Notes:

- a. This command can interrupt the formation of a consistency group. If this command does not complete, an error code is issued. If this occurs, examine the error code and follow your local procedures for problem determination. In most cases, if you correct the error, you can successfully reissue the **rmgmir** command. If, however, the reissued **rmgmir** command fails and Global Mirror processing must be stopped, reissue the **rmgmir** command with the **-force** parameter.
- b. You cannot use the **rmgmir** command to stop a script that involves Global Mirror processing. The only way to stop a script is to press the **Ctrl C** keys on your keyboard. This action stops the DS CLI session. However, it does not stop the microcode that is processing Global Mirror transactions. To stop the microcode processing, you must log back into the DS CLI session and issue the **rmgmir** command.
- 4. Enter Y to confirm that you want to stop Global Mirror processing for the specified session.

The message is displayed as follows:

Are you sure you want to stop Session ID (xx)? [y/n]: Y Global Mirror for Session ID 01 successfully stopped.

5. Issue the **mkgmir** command with your updated master and subordinate LSS changes to start Global Mirror processing. Enter the **mkgmir** command at the dscli command prompt with the following parameters and variables:

dscli> mkgmir -dev storage_image_ID -lss ID -session session_ID -cginterval seconds -coordinate milliseconds -drain seconds Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID

Example

dscli> mkgmir -dev IBM.2107-75FA120 -lss 10 -session 01 -cginterval 0 -coordinate 40 -drain 35 IBM.2107-75FA120/00:IBM.2107-75FA150/00

Notes:

a. You can change your mind and decide not to change any of the topology values. However, you must still issue the **mkgmir** command to resume Global Mirror processing after you have stopped

the processing. Because the **resumegmir** command is used only with the **pausegmir** command, you cannot issue the **resumegmir** command to restart the processing.

b. When you stopped the Global Mirror session, the values for the tuning parameters become invalid. You must specify values for these parameters (**-cginterval**, **-coordinate**, **-drain**) when you restart your Global Mirror session.

Viewing a Global Mirror session

Complete this task to view the associated properties of the Global Mirror session or information about the volumes in each session.

Issue the **lssession** command to display Global Mirror session information regarding the volumes in each session.

Complete the following step to view the associated properties of the Global Mirror session or information about Global Mirror session failures. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with values declared for the variables.

Issue the **lssession** command to display Global Mirror session information regarding the volumes in each session. Enter the **lssession** command at the dscli command prompt with the following parameters and variables:

dscli> lssession -dev storage_image_ID -l LSS_ID

Example

dscli> lssession -dev IBM.2107-75FA120 -l 01

Note:

- Use the -1 parameter if you want to see a detailed report. The report provides the following information:
 - State of the session status. For example, whether the consistency group of the session is in progress or the increment process is in progress.
 - The status of each volume in the session.
 - Whether the first cycle of the volume in the global mirror relationship has ended. This value is displayed as true or false.
- Use the **-s** parameter if you only want to see the volumes (with no details) that are associated with each session within the LSS.

Querying Global Mirror processing

Complete this task to display detailed properties and performance metrics for a Global Mirror session.

Issue the **showgmir** command to display a detailed report about the current Global Mirror operations and to request that the report include the detailed performance metrics.

Complete the following step to view the detailed properties and performance metrics for a Global Mirror session. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **showgmir** command to display a detailed report about the current Global Mirror operations. Or, you can request a report that displays the performance metrics associated with the current Global Mirror operations. Enter the **showgmir** command at the dscli command prompt with the following parameters and variables:

For transaction details, enter: dscli> showgmir -dev *storage_image_ID LSS_ID* For performance metrics, enter: dscli> showgmir *storage_image_ID* -metrics *LSS_ID* **Examples**

dscli> showgmir -dev IBM.2107-75FA120 10 dscli> showgmir -dev IBM.2107-75FA120 -metrics 10

Pausing Global Mirror processing

Complete this task to pause Global Mirror processing.

Use the **pausegmir** command to pause Global Mirror processing. This pause function allows you to temporarily suspend Global Mirror processing attempts to form consistency groups.

There are 2 primary reasons to pause Global Mirror processing:

- You must repair a part of the Global Mirror infrastructure, such as:
 - Global Copy volume pairs
 - FlashCopy pairs
 - Storage control unit values
- You must make modifications to the Global Mirror tuning parameters

Complete the following steps to pause Global Mirror processing. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **pausegmir** command to pause Global Mirror processing. Enter the **pausegmir** command at the dscli command prompt using the following parameters and variables:

dscli> pausegmir -dev storage_image_ID -lss LSS_ID -session session_ID Master Control Path LSS ID:Subordinate Control Path LSS ID

Example

dscli> pausegmir -dev IBM.2107-75FA120 -lss 10 -session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00

2. Use the **showgmir** command to verify that your changes are correct. When you are ready to resume Global Mirror processing, issue the **resumegmir** command to continue with the Global Mirror processing. Do not issue the start (**mkgmir**) command to start Global Mirror processing.

Resuming Global Mirror processing

Complete this task to resume Global Mirror processing after you have paused Global Mirror processing.

Note: If you have issued a **pausegmir** command to pause Global Mirror processing, issue the **resumegmir** command to continue Global Mirror processing.

Use the **resumegmir** command to change your Global Mirror tuning parameters and continue Global Mirror processing. When you change the Global Mirror tuning parameters, you must include values for all three parameters (consistency group interval time, coordination interval time, and drain time). You cannot submit a value for just one parameter, even if the two other parameters do not need to be changed.

Complete the following steps to resume Global Mirror processing. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **resumegmir** command to continue Global Mirror processing after you have paused Global Mirror processing. Enter the **resumegmir** command at the dscli command prompt using the following parameters and variables:

dscli> resumegmir -dev storage_image_ID -lss LSS_ID -session session_ID Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID Example dscli> resumegmir -dev IBM.2107-75FA120 -lss 10 -session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00

Note: If you are making changes to your tuning parameters, your command looks like the following example:

```
dscli> resumegmir -dev IBM.2107-75FA120 -lss 10 -cginterval 5
-coordinate 50 -drain 30 -session 01
IBM.2107-75FA120/00:IBM.2107-75FA150/00
```

In this example the **-cginterval** parameter was changed while the **-coordinate** and **-drain** parameters maintained their DS CLI default values. However, because the **-cginterval** parameter was changed, all the parameters and their corresponding values must be listed in your command. Otherwise, the command fails.

Starting Global Mirror processing

Complete this task to start Global Mirror processing.

The volume relationships (paths, pairs, and FlashCopy) plus the creation of a Global Mirror session must be complete before Global Mirror processing can start.

Use the **mkgmir** command to start Global Mirror processing.

Complete the following step to start Global Mirror processing. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Issue the **mkgmir** command to start Global Mirror processing. Enter the **mkgmir** command at the dscli command prompt using the following parameters and variables:

dscli> mkgmir -dev storage_image_ID -lss LSS_ID -cginterval seconds -coordinate milliseconds -drain seconds -session session_ID Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID

Example

dscli> mkgmir _dev IBM.2107-75FA120 -lss 10 -cginterval 0 -coordinate 50 -drain 30 -session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00

Note: Issuing the **mkgmir** command requires that you specify the tuning parameters. The values for the tuning parameters are not retained when you end Global Mirror processing. So, in the case where you must change the Global Mirror topology parameters, resubmit the tuning parameters when you restart Global Mirror processing.

Ending Global Mirror processing (script mode)

Complete this task to end Global Mirror processing that is being controlled through a script. Complete the steps in this task only if there is no alternative. For example a disaster has occurred and you must immediately stop all processing.

There is no way to pause a script. The only way to stop a script is to press the **Ctrl C** keys, which stops your DS CLI session. It is likely that this action might cause some transactions to remain partly completed and others undone.

Pressing **Ctrl C** does not stop Global Mirror processing which is controlled through the microcode. To stop the microcode processing of Global Mirror operations, you must log back into a DS CLI session and issue the **rmgmir** command.

This task does not provide the detailed instructions for recovery. The recovery instructions are described in the *Recovering when a disaster strikes* task.

Complete the following step when you want to end Global Mirror processing that is using a script.

- 1. Press the Ctrl C keys to immediately end the DS CLI session.
- 2. Log back into a DS CLI session and enter the **rmgmir** command to stop the microcode processing of the Global Mirror operations.
- 3. Proceed with the steps that are described in the *Recovering when a disaster strikes* task.

Ending Global Mirror processing (no script)

Complete this task to end Global Mirror processing that is not being controlled through a script.

To use this task your Global Mirror processing cannot be controlled through a script.

You can use this task when you must end Global Mirror processing to change the topology of a Global Mirror session or when you have time (because of a rolling disaster) to end processing even though a disaster has occurred. The **rmgmir** command is used to end Global Mirror processing.

Note: This command might interrupt the formation of a consistency group. If, due to failures, this command cannot complete an error code is issued. If this occurs, examine the error code and follow your local procedures for problem determination. In most cases, correcting the error and reissuing the **rmgmir** command is successful. However, if reissuing the **rmgmir** command fails and Global Mirror absolutely must be ended, the **rmgmir** command can be reissued with the **-force** parameter.

Complete the following steps to end Global Mirror processing. The example commands displayed in this task are shown in two formats. The first format shows the type of information the command requires. The second format provides the command with declared values for the variables.

1. Issue the **rmgmir** command to end Global Mirror processing. Enter the **rmgmir** command at the dscli command prompt using the following parameters and variables:

dscli> rmgmir -dev storage_image_ID -lss LSS_ID -session session_ID Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID

Example

dscli> rmgmir -dev IBM.2107-75FA120 -lss 10 -session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00

2. Enter Y in response to each message that requests that you confirm that you want the specified session stopped. A message like the following example appears when you process the **rmgmir** command.

Are you sure you want to close Session ID 01? [y/n]: Y Global Mirror for Session ID 01 closed successfully.

Setting up the Global Mirror Environment

This task lists the high-level steps that you must complete to set up the Global Mirror environment to allow Global Mirror processing.

Each of these steps must be completed in the order in which they are shown before you can move onto the next step.

- 1. Create Fibre Channel paths between all Global Mirror source and target pairs and between the Master and subordinate storage units. See "Creating Fibre Channel paths (Global Mirror setup)" on page 592 for additional steps.
- 2. Create Global Copy pairs from the local storage units to the remote storage units. See "Creating Global Copy pairs (Global Mirror setup)" on page 593 for additional information.

- **3**. Create FlashCopy relationships at the remote site between the Global Copy secondary volumes and the FlashCopy target volumes. See "Creating FlashCopy relationships (Global Mirror setup)" on page 594 for additional information.
- 4. Create the Global Mirror session. See "Creating the Global Mirror session" on page 595 for additional information.

Creating Fibre Channel paths (Global Mirror setup)

Complete this task to create Fibre Channel paths between all Global Mirror source and target pairs and between the Master and subordinate storage units. This task is the first step in setting up your Global Mirror environment.

Create paths so that the logical subsystems (LSSs) are associated with each other. Copy services I/O pass through these ports. It is preferred that they are not the same ports that are used for host I/O. This setup ensures that there is enough capacity for the data transfer.

Complete the following steps to create Fibre Channel paths between all Global Mirror source and target pairs and between the Master and subordinate storage units. The example commands displayed in this task are shown in two formats. The first format shows the type of information the command requires. The second format provides the command with declared values for the variables.

 Obtain the worldwide node name of the secondary storage unit. This information is needed when you do the next step. Enter the lssi or showsi at the dscli command prompt as follows: dscli> lssi -1

This input is the entire command. No additional variables are needed.

The **showsi** command does contain a variable and a command parameter: dscli> showsi *storage_image_id* -fullid

Example dscli> showsi -fullid IBM.2107-75FA120

Notes:

- a. Use the **lssi** command if you want to display a list of all the storage image instances for a storage-complex and status information for each storage image in the list.
- b. Use the **showsi** command if you want to display the detailed properties of a specific storage unit.
- c. Use the **-fullid** parameter with the **showsi** command to display fully qualified IDs, which include the storage image ID, for every ID that is displayed in the command output.
- d. Record the worldwide node name for the secondary (target) storage unit so that it can be used when you issue the **mkpprcpath** command.
- 2. Issue the **lsavailpprcport** command to display a list of Fibre Channel I/O ports that can be defined as remote mirror and copy paths. Enter the **lsavailpprcport** command at the dscli command prompt with the parameters and variables shown as follows:

dscli> lsavailpprcport -dev storage_image_ID -remotedev storage_image_ID -remotewwnn wwnn source_LSS_ID:target_LSS_ID

Example

```
dscli> lsavailpprcport -l -dev IBM.2107-75FA120
-remotedev IBM.2107-75FA150 -remotewwnn 12341234000A000F 01:01
```

Notes:

- a. The -remotedev parameter specifies the ID of the secondary storage unit.
- b. The **-remotewwnn** parameter must specify the worldwide node name of the secondary storage unit. If you make a mistake and specify the worldwide node name of the primary storage unit, the command fails.

c. The shortened version of the source_LSS_ID:target_LSS_ID parameter is shown (value = 01:01) because the example uses the fully qualified -devstorage_image_ID and -remotedevstorage_image_ID parameters. If the fully qualified -dev and -remotedev parameters are not used, you must use the fully qualified source_LSS_ID:target_LSS_ID value. For example: IBM.2107-75FA120/01:IBM.2107-75FA150/01.

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

3. Issue the **mkpprcpath** command to create the Fibre Channel paths between all Global Mirror source and target pairs and between the Master and subordinate storage units. Enter the **mkpprcpath** command at the dscli command prompt with the following parameters and variables as follows:

dscli> mkpprcpath -dev storage_image_ID -remotedev storage_image_ID -remotewwnn wwnn -srclss source_LSS_ID -tgtlss target_LSS_ID source_port_ID:target_port_ID

Example

```
dscli> mkpprcpath -dev IBM.2107-75FA120 -remotedev
IBM.2107-75FA150 -remotewwnn 12341234000A000F
-srclss IBM.2107-75FA120/00 -tgtlss IBM.2107-75FA150/01 I1A10:I2A20
```

Notes:

- a. The -remotedev parameter specifies the ID of the secondary storage unit.
- b. The **-remotewwnn** parameter must specify the worldwide node name of the secondary storage unit. If you make a mistake and specify the worldwide node name of the primary storage unit, the command fails.
- c. The shortened version of the -srclss parameter is shown (value = 00) because the example uses the fully qualified -dev storage_image_ID parameter. If the fully qualified -dev parameter was not used, you must use the fully qualified -srclsssource_LSS_ID value. For example: -srclss IBM.2107-75FA120/00.
- d. The shortened version of the -tgtlss parameter is shown (value = 01) because the example uses the fully qualified -dev storage_image_ID parameter. If the fully qualified -dev parameter was not used, you must use the fully qualified -tgtlsstarget_LSS_ID value. For example: -tgtlss IBM.2107-75FA120/01.
- e. The shortened version of the *source_port_ID:target_port_ID* parameter is shown (value = I1A10:I2A20) because the example uses the fully qualified -devstorage_image_ID and -remotedevstorage_image_ID parameters. If the fully qualified -dev and -remotedev parameters were not used, you must use the fully qualified source_port_ID:target_port_ID value. For example: IBM.2107-75FA120/I1A10:IBM.2107-75FA150/I2A20.

The fully qualified *source_port_ID:target_port_ID* parameter is positional on your command line. It must be placed after the -tgtlss parameter and value. For example:

dscli> mkpprcpath -srclss 00 -tgtlss 01 IBM.2107-75FA120/I1A10:IBM.2107-75FA150/I2A20

Creating Global Copy pairs (Global Mirror setup)

Complete this task to create Global Copy pairs from the local storage units to the remote storage units. This is the second step in setting up your Global Mirror environment.

Ensure that you have already created the Fibre Channel paths between all Global Mirror source and target pairs and between the master and subordinate storage units.

The purpose of this step is to create a relationship between a source volume and a target volume.

Complete the following steps to create Global Copy pairs from the local storage units to the remote storage units. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

1. Issue the **mkpprc** command to create a Global Copy pair from the local storage unit to the remote storage unit and to create a relationship between the associated source volume and target volume. Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID -type
gcp SourceVolumeID:TargetVolumeID
```

Example

```
dscli> mkpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -type gcp 0100:0100
```

Notes:

- a. The **-remotedev** parameter specifies the ID of the secondary storage unit.
- b. The **-type** *gcp* parameter specifies that one or more remote mirror and copy Global Copy volume relationships be established. Global Copy maintains the remote mirror and copy relationship in a nonsynchronous manner.

A confirmation message is issued for each successful Global Copy volume association that is created.

2. Issue the **1spprc** command to check the status information for each remote mirror and copy volume relationship in the list. Enter the **1spprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID
    -l SourceVolumeID:TargetVolumeID
```

Example

dscli> lspprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -1 0100:0100

Use the -1 parameter to provide a more detailed report about the Global Copy volume relationships.

Wait until the Global Copy pair process has completed its first pass before you begin creating the FlashCopy relationships.

Note: Global Copy source volumes are not in the active Global Mirror session until the volumes have been added to the session and the session has started.

Creating FlashCopy relationships (Global Mirror setup)

Complete this task to create FlashCopy relationships at the remote site between the Global Copy secondary volumes and the FlashCopy target volumes. This is the third step in setting up your Global Mirror environment.

The following tasks must be completed before you can initiate this step:

- Ensure that the Fibre Channel paths between all Global Mirror source and target pairs and between the master and subordinate storage units are created.
- Ensure that the Global Copy pairs are created between the local storage units and the remote storage units.

The purpose of this task is to create a FlashCopy target for the Global Mirror pairs.

Complete the following steps to create FlashCopy relationships at the remote site between the Global Copy secondary volumes and the FlashCopy target volumes. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: Enter the commands that are described in the steps either for a DS8000 model or for a DS6000 model. The storage image ID for the DS6000 model is different.

1. Enter the **mkflash** command to create FlashCopy relationships at the remote site between the Global Copy secondary volumes and the FlashCopy target volumes.

```
dscli> mkflash -dev storage_image_ID -tgtinhibit -persist -record -nocp
sourcevolumeID:targetvolumeID
```

Example

```
dscli> mkflash -dev IBM.2107-75FA150 -tgtinhibit -record -persist -nocp 0001:0004
```

Notes:

- a. Specify the storage image ID of the secondary storage unit for the **-dev** *storage_image_ID* parameter. If the management console has an IP connection to the specified remote site, the command works. If the IP connection is not established, enter the **mkremoteflash** command with all the same parameters as displayed in the example.
- b. Specify the **-tgtinhibit** parameter to prevent writes to the target volume.
- c. Specify the **-record** parameter to activate the change recording function on the volume pair.
- d. Specify the **-persist** parameter to retain the FlashCopy relationship after the background copy completes.
- e. Specify the -nocp parameter to prevent creating a background copy.

A confirmation message is generated for each successful FlashCopy pair that is created.

2. Enter the **lsflash** command to check the status information for each FlashCopy relationship at the remote site.

dscli> lsflash -dev storage_image_ID -l SourceVolumeID:TargetVolumeID

Example

dscli> lsflash -dev IBM.2107-75FA150 -1 0100:0100

Use the -1 parameter to provide a more detailed report about the FlashCopy relationships.

Note: If you used the **mkremoteflash** command, you must enter the **lsremoteflash** command to run a status check.

Creating the Global Mirror session

Complete this task to create a Global Mirror session. This is the fourth step in setting up your Global Mirror environment. After you complete this step you are ready to start Global Mirror processing.

The following tasks must be completed before you can proceed with this step.

- Ensure that the Fibre Channel paths between all Global Mirror source and target pairs and between the master and subordinate storage units have been created.
- Ensure that the Global Copy pairs have been created between the local storage units and the remote storage units.
- Ensure that the FlashCopy relationships at the remote site between the Global Copy secondary volumes and the FlashCopy target volumes have been created.

The purpose of this step is to create a container that associates volumes with a Global Mirror session.

Complete the following steps to create the Global Mirror session. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **mksession** command to create the Global Mirror session. Enter the **mksession** command at the dscli command prompt using the following parameters and variables:

```
dscli>mksession -dev storage_image_ID -lss LSS_ID -volume volume_ID Session_ID
```

Example

dscli> mksession -dev IBM.2107-75FA120 -lss 10 -volume 0800-0810 01

2. Repeat Step 1 for each LSS.

You must make a session for each LSS. However, you can associate each LSS with the same session. For example: You have LSS 08 and it contains volumes 0800 - 0810. You create a session and assign it to session 08. You also have LSS 09 and it contains volumes 0900 - 0910. You create a session and assign it to session 08. When you start Global Mirror processing, the volumes for LSS 8 and LSS 9 are processed in the same session (session 08).

Removing a Global Mirror environment

This task lists the high-level steps that you must complete to remove the Global Mirror environment from your system.

The removal of a Global Mirror environment is prompted by circumstances like the following cases:

- You decide that you want to run with a different configuration.
- You were running a test and the volumes you were using you never want to use for an asynchronous pair again.

Each of the following steps must be completed in sequence before you can proceed to the next step.

- 1. Remove all Global Copy primary volumes from the Global Mirror sessions. See "Removing volumes from a session (Global Mirror)" for additional information.
- 2. End the Global Mirror sessions.

See "Ending a Global Mirror session" on page 597 for additional information.

- **3**. Withdraw all FlashCopy relationships between the B and C volumes. See "Removing FlashCopy relationships" on page 598 for additional information.
- 4. Remove the Global Copy pair relationships. See "Removing the Global Copy pair relationship" on page 599 for additional information.
- 5. Remove the remote mirror and copy paths between the local site and the remote site. "Removing the Fibre Channel paths" on page 600 for additional information.

Removing volumes from a session (Global Mirror)

Complete this task to remove volumes from a Global Mirror session. This task is also the first step in removing the Global Mirror environment from your system.

Ensure that you have closed the Global Mirror sessions associated with the Global Mirror environment if you are processing this task as part of removing the Global Mirror environment from your system.

This task covers two circumstances

- · removing volumes from a session without removing the session from Global Mirror processing
- removing all volumes from each session as part of the steps for removing the Global Mirror environment from your system.

You can remove Global Copy primary volumes from a Global Mirror session at any time after the Global Mirror session has started without stopping the session.

If you have many volumes that you want to remove from a Global Mirror session, you might consider removing them from the session in stages. This lessens the impact on your processing. If you intend to use the volumes in a different configuration, you must remove the pair and path associations. Removing a volume from a Global Mirror session does not remove the pair and path associations.

Complete the following steps to remove volumes from a Global Mirror session. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **lssession** command to query the status of all volumes that are associated with the sessions of a specified logical subsystem. Enter the **lssession** command at the dscli command prompt with the following parameters and variables:

dscli> lssession -dev storage_image_ID -l LSS_ID

dscli> lssession -dev IBM.2107-75FA120 -1 01

When you use the **-1** parameter, a detailed report displays a list of Global Mirror sessions for the specified logical subsystem (LSS) and information about the volumes of each session in the list.

- 2. Analyze the report and determine which volumes that you want to remove.
- **3**. Issue the **chsession** command to remove the specified volumes from the specified Global Mirror session. Enter the **chsession** command at the dscli command prompt with the following parameters and variables:

dscli> chsession -dev storage_image_ID -lss LSS_ID -action remove -volume volume_ID session_ID

Example

dscli> chsession -dev IBM.2107-75FA120 -lss 10 -action remove -volume 0100-010F,0180-018F,0120 01

A confirmation message indicates that the session has been modified successfully.

Note: 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.

After the volumes have been removed from the Global Mirror session, you must end the volume associations for the removed volumes (FlashCopy, Global Copy pair, and remote mirror and copy path) if you plan to use the volumes in a different configuration.

Ending a Global Mirror session

Complete this task to end a Global Mirror session. In addition to ending a single session, this task is also the second step that you use when you remove the Global Mirror environment from your system.

Each session that you have created for Global Mirror processing must be ended individually. You cannot designate that a range of sessions be ended. Ending a session does not remove the volume or path associations. Each of these must be removed in their own way.

Complete the following steps to end a Global Mirror session. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **Issession** command to obtain a list of all sessions that are associated with the specified logical subsystem. Enter the **Issession** command at the dscli command prompt with the following parameters and variables:

```
dscli> lssession -dev storage_image_ID -s LSS_ID
```

Example

dscli> lssession -dev IBM.2107-75FA120 -s 01

Note: For the circumstance described in this task, it is better to issue the **-s** parameter. The **-s** parameter creates a report with the following three items of information:

- LSSID
- Session number
- Volume numbers
- 2. Print the report or record the session numbers that need to be ended.
- **3**. Issue the **rmsession** command to end the specified session. Enter the **rmsession** command at the dscli command prompt with the following parameters and variables:

dscli> rmsession -dev storage_image_ID -lss ID session_ID

Example

dscli> rmsession -dev IBM.2107-75FA120 -lss 10 01

4. Enter **Y** to respond to the message that requests that you confirm that you want to end the specified session. A message like the following example is displayed when you process the **rmsession** command.

Are you sure you want to close Session ID 01? y/n Y Global Mirror Session ID 01 closed successfully.

5. Repeat Step 3 for each session that you want to end.

Removing FlashCopy relationships

Complete this task to remove FlashCopy relationships that exist with the volumes that are part of your Global Mirror environment. In addition, this task is the third step in removing the Global Mirror environment from your system.

Ensure that the following tasks have been completed before this step when you are removing your Global Mirror environment:

- Remove the volumes that are associated with each Global Mirror session you have closed.
- End the Global Mirror sessions that are part of your Global Mirror environment.

If you attempt to remove the FlashCopy relationships (as part of removing the Global Mirror environment) before doing the first two tasks, this task fails.

This task step only applies as part of the removal of your Global Mirror environment.

Note: Make certain that you no longer need your consistent set of data before you withdraw your relationships. If you still need your consistent data, initiate a data backup before proceeding with this task.

Complete the following steps to remove FlashCopy relationships that existed with the volumes that were part of your Global Mirror environment. The example commands in this task are shown in two formats. The first format shows the type of information the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **lsflash** command to check the status information for each FlashCopy relationship at the remote site. A detailed report (when you use the **-l** parameter) is displayed for each FlashCopy relationship. Enter the **lsflash** command at the dscli command prompt with the parameters and variables shown as follows:

dscli> lsflash -dev storage_image_ID -l SourceVolumeID:TargetVolumeID.

dscli> lsflash -dev IBM.2107-75FA150 -l 0100:0100

Use the -l parameter to provide a more detailed report about the FlashCopy relationships.

Note: If you originally used the **mkremoteflash** command to create your FlashCopy relationships, you must enter the **lsremoteflash** command to initiate a status check.

- 2. Analyze the list of volumes that have been part of your Global Mirror environment and ensure that these are the volumes from which the FlashCopy relationship must be removed.
- **3**. Issue the **rmflash** command to remove the FlashCopy volume relationships. Enter the **rmflash** command at the dscli command prompt with the parameters and variables shown as follows:

dscli> rmflash -dev storage_image_ID SourceVolumeID:TargetVolumeID

Example

dscli> rmflash -dev IBM.2107-75FA120 0001:0004 0003:00FF 0008:000C

Note:

- The example shows the use of multiple FlashCopy pair IDs. Ensure that you separate multiple FlashCopy pair IDs with spaces.
- If you used the **mkremoteflash** command to create your FlashCopy relationships, you must enter the **rmremoteflash** command to remove the FlashCopy relationships.
- 4. Enter **Y** in response to each message that requests that you confirm that you want the specified FlashCopy pair removed. A message like the following example appears for each FlashCopy pair being removed when you process the **rmflash** command.

Are you sure you want to remove the FlashCopy pair 0001:0004? [y/n]: Y

FlashCopy pair 0001:0004 successfully removed.

Removing the Global Copy pair relationship

Complete this task to remove Global Copy relationships that existed with the volumes that were part of your Global Mirror environment. This task is the fourth step in removing the Global Mirror environment from your system.

Ensure that you have completed the following tasks before you initiate this task; otherwise, this task fails:

- Remove the volumes that are associated with each Global Mirror session that you have closed.
- End the Global Mirror sessions that are part of your Global Mirror environment.
- Remove FlashCopy relationships that exist with the volumes that were part of your Global Mirror environment.

The purpose of this task is to remove the Global Copy relationships for each pair of source volumes on your primary site and the target volumes on your secondary site. There might be several LSSs that are involved and the Global Copy relationships must be removed within each LSS. This requires that you issue the **rmpprc** command for each LSS until all relationships are removed.

Complete the following steps to remove Global Copy pair relationships that exist with the volumes that were part of your Global Mirror environment. The example commands in this task are shown in two formats. The first format shows the type of information the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **lspprc** command to check the status information for each Global Copy volume relationship in the list. Enter the **lspprc** command at the dscli command prompt with the parameters and variables shown as follows:

dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID -l SourceVolumeID:TargetVolumeID.

Example

dscli> lspprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -1 0100:0100

Use the -l parameter to provide a more detailed report about the Global Copy volume relationships.

2. Analyze the list of volumes that have been part of your Global Mirror environment, and ensure that these are the volumes from which the Global Copy relationships must be removed.

3. Issue the **rmpprc** command to remove the Global Copy volume relationships. Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -dev *storage_image_ID* -remotedev *storage_image_ID* SourceVolumeID:TargetVolumeID **Example**

dscli> rmpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 0100:0100

4. Enter **Y** in response to each message to confirm that you want to remove the specified Global Copy pair. A message like the following lines is displayed for each Global Copy pair that is being removed when you process the **rmflash** command.

Are you sure you want to remove PPRC pair 0100:0100? [y/n]: Y

Remote Mirror and Copy pair IBM.2107-75FA120/0100:0100 successfully removed.

5. Repeat Steps 3 and 4 for all the volumes that have Global Copy relationships in the LSSs that were part of your Global Mirror environment.

Removing the Fibre Channel paths

Complete this task to remove the Fibre Channel paths between all Global Mirror source and target pairs and between the master and subordinate storage units. This is the fifth and final step in removing the Global Mirror environment from your system.

Ensure that you have completed the following tasks before you initiate this task; otherwise, this task fails:

- Remove the volumes that are associated with each Global Mirror session that you have closed.
- Close the Global Mirror sessions that are part of your Global Mirror environment.
- Remove FlashCopy relationships that exist with the volumes that were part of your Global Mirror environment.
- Remove the Global Copy relationships that exist with the volumes that were part of your Global Mirror environment.

Repeat the process to remove the Fibre Channel paths for each data path and for each control path. Data paths consists of source LSSs on the storage images at the primary site and target LSSs on the storage images at the secondary site. Control paths consist of the master storage image LSS and the subordinate storage image LSSs.

Complete the following steps to remove the Fibre Channel paths between all Global Mirror source and target pairs and between the master and subordinate storage units. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **lspprcpath** command to display a list of existing remote mirror and copy path definitions. Enter the **lspprcpath** command at the dscli command prompt with the following parameters and variables:

dscli> lspprcpath -dev storage_image_ID source_LSS_ID

Example

dscli> lspprcpath -l -dev IBM.2107-75FA120 01

Note: The report that is displayed from this command provides the worldwide node name that is used with the **rmpprcpath** command.

2. Record the path information to use when you issue the **rmpprcpath** command.

3. Issue the **rmpprcpath** command to remove the Fibre Channel paths between all Global Mirror source and target pairs and between the master and subordinate storage units. Do this for each path that must be removed. Enter the **rmpprcpath** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprcpath -dev storage_image_ID -remotedev storage_image_ID -remotewwnn wwnn source_LSS_ID:target_LSS_ID

Example

dscli> rmpprcpath -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -remotewwnn 12341234000A000F 01:01

4. Enter **Y** in response to each message to confirm that you want to remove the specified remote mirror and copy path. A message like the following lines is displayed for each remote mirror and copy path that is being removed.

Are you sure you want to delete the Remote Mirror and Copy path (whatever was designated)? [y/n]: Y

Remote Mirror and Copy path (designated in the command) successfully deleted.

Note: Use of the **-quiet** parameter with the **rmpprcpath** command turns off the confirmation question for each path that is being removed.

5. Repeat Step 3 for all the remote mirror and copy paths per LSS that were part of your Global Mirror environment.

After this task is complete, you can create your new Global Mirror environment and configuration.

Chapter 8. Recovering from a disaster using Global Mirror

Use this information to complete the high-level steps that must be done to recover from a disaster using Global Mirror processing.

A failure at the local or primary site stops all I/O to and from the local storage server. The local server cannot communicate with the remote sites. This might impact the formation of consistency groups, because the entire process is managed and controlled by the master storage server, which is the primary storage server.

Your initial goal is to swap operations between the local and remote sites and then restart the applications. This requires that you make available a set of consistent volumes at the remote site, before the application can restart at the remote site.

When the local site is operational again, you want to return processing to the local site. Before you can return processing to the local site, you must apply changes from the remote site to the local site. These changes are the transactions that occurred after you started failover processing to the remote site.

The following considerations can help you determine where transactions are being processed:

- The **local site** contains A volumes (the source volume), which are copied to the recovery site using Global Copy
- The **recovery (or remote) site** contains B volumes (the target volume and FlashCopy source volume) and C volumes (the FlashCopy target volume)
- A storage unit at the local site is designated as the Global Mirror master and all other local (or production) storage units are designated as subordinate storage units. The master storage unit sends commands to its subordinate storage units. These subordinates work together to create a consistency group and to communicate the FlashCopy commands to the recovery (or remote) site. All status is relayed back to the Global Mirror master.

To recover from a disaster, you must complete the following high-level tasks using the Global Mirror function and the DS CLI commands:

- End Global Mirror processing when a disaster occurs.
 See "Ending Global Mirror processing when a disaster occurs" on page 604 for additional substeps.
- 2. Check the status of the current processing for Global Mirror transactions. See "Checking Global Mirror transaction status in a disaster situation" on page 604 for additional substeps.
- **3**. Initiate the failover process of A volumes to B volumes. See *"Initiating failover processing for B volumes to A volumes" on page 605* for additional substeps.
- 4. Analyze the consistency group status. See "Analyzing and validating the consistency group state" on page 606 for additional substeps.
- 5. Use the **revertflash** command to correct FlashCopy relationships. See "Using the revertflash command to correct FlashCopy relationships" on page 608 for additional substeps.
- 6. Use the **commitflash** command to correct FlashCopy relationships. See "Using the commitflash command to correct FlashCopy relationships" on page 609for additional substeps.
- 7. Initiate the fast reverse restore process. See "Using fast reverse restore processing to create consistency" on page 610 for additional substeps.
- 8. Wait for the background copy to complete. See "Waiting for the background copy to complete" on page 611 for additional substeps.
- **9**. Reestablish the FlashCopy relationships, B volumes to C volumes. See *"Reestablishing the FlashCopy relationships between B volumes and C volumes" on page 611* for additional substeps.

- **10**. Prepare to reinstate production at the local site. See "*Preparing to reinstate production at the local site*" *on page 612* for additional substeps.
- 11. Resynchronize the volumes. See "Resynchronizing the volumes" on page 613 for additional substeps.
- 12. Query for first pass and drain time out-of-synch zero value and quiesce your system. See "Querying, quiescing, and re-querying" on page 614 for additional substeps.
- **13**. Reestablish the remote mirror and copy paths, A site to B site. See *"Reestablishing remote mirror and copy paths (site A to site B)" on page 615* for additional substeps.
- 14. Run Global Copy failover processing to the A volumes. See *"Running Global Copy failover processing to the A volumes" on page 617* for additional substeps.
- **15**. Run Global Copy failback processing for the A volumes. See *"Running Global Copy failback processing for the A volumes" on page 617* for additional substeps.
- **16**. Resume Global Mirror processing at site A. See *"Resuming Global Mirror processing at site A" on page 618* for additional substeps.

Ending Global Mirror processing when a disaster occurs

Complete this task to end Global Mirror processing when a disaster occurs. This is the first step in the Global Mirror disaster recovery process.

Depending on the state of the local Global Mirror storage server, you might have an opportunity to end Global Mirror processing before you can initiate the rest of the recovery steps. A disaster can affect your local server and your choices for ending Global Mirror processing in one of the following ways:

• Your site experiences a rolling disaster and your Global Mirror processing is not being done through a DS CLI script. This circumstance allows you time to issue a pause command, followed by a query command, and then an end Global Mirror processing command.

Note: If the query displays a status of Fatal or a null (-), you must analyze and correct your consistency groups during the recovery process. If the query displays a status of Paused, your consistency groups are formed before you end Global Mirror processing.

• Your site experiences a rolling disaster and your Global Mirror processing is being done through a DS CLI script. You can take a chance that you have enough time to allow the script to process to the end. However, it is likely that you must end Global Mirror processing by pressing the **CTRL C** buttons on your keyboard because there is no pause feature when running a script.

Because this is a rolling disaster, you might have time to log back into your DS CLI session on the local server and query the status of the Global Mirror processing before it was ended.

• Your sites' local server is affected by the disaster immediately and you have no time to end Global Mirror processing or issue a status query.

If possible, you want to end Global Mirror processing. Complete the following step to end Global Mirror processing.

Issue the **rmgmir** command to end Global Mirror processing or press the **CTRL C** buttons when you are using a DS CLI script and then issue the **rmgmir** command.

After Global Mirror processing has been ended, you are ready for the next step in the Global Mirror failover recovery process.

Checking Global Mirror transaction status in a disaster situation

Complete this task to check the status of your Global Mirror transactions before processing ends as a result of a disaster. Normally, this is the second step in the Global Mirror failover recovery process because in a rolling disaster, your primary server is still operational. This might be the first step if your primary server is not operational.

You must end Global Mirror processing before you can initiate this task.

You must determine the status of the Global Mirror processing before processing ends. Some transactions might be half completed while others are not yet started. Querying the status of your transactions provides a basis for planning which tasks must be done next. Your situation can be one of the following conditions:

- When Global Mirror processing ends, the formation of a consistency group is in progress and the FlashCopy state between the B and the C volumes in the remote storage server is not the same for all relationships.
- Some FlashCopy pairs might have completed the FlashCopy phase of Global Mirror processing to form a new consistency group and might have already committed the changes.
- Some FlashCopy pairs might not have completed and are in the middle of processing to form their consistent copy and remain in a revertible state.
- There is no master server that controls and coordinates the processing that might continue for a brief period at the remote site.

Complete the following steps to obtain an initial status of your transactions.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

- 1. If your primary server is still operational, issue the **lssession** and **lspprc** commands to obtain reports that allow you to determine the status of your Global Mirror transactions. If your primary server is not available, go to the next step.
- 2. Gain access to your remote server and navigate to the directory where you installed the DS CLI.
- **3**. Log in to a DS CLI session.
- 4. Issue the **1spprc** command to provide a report that allows you to determine the status of your Global Mirror transactions.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID -l SourceVolumeID:TargetVolumeID

Example

dscli> lspprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -1 0100-0103:0100-0103 0200:0200 0300-0303:0300-0303

Note: The report displays your Global Copy pairs with a suspended state.

After you obtain your reports, you are ready for the next step in the Global Mirror disaster recovery process, which is to issue the failover command.

Initiating failover processing for B volumes to A volumes

Complete this task to run failover processing for B volumes to A volumes so that the B volumes become the primary volumes and A volumes become the secondary volumes. The Global Copy state of the B volumes changes from secondary to primary and suspended. This is the next step after you obtain a status of your Global Mirror transactions in the Global Mirror recovery failover process.

Failover processing cannot be run unless following tasks are complete:

- Global Mirror processing complete
- Reports that help you determine the status of your Global Mirror transactions before the disaster occurred are available

Recovery failover processing on the Global Copy volume pair ends the A-to-B volumes extended distance relationship and instead creates a B-to-A volume Global Copy relationship. Therefore, failover processing to the Global Copy secondary volume turns the secondary volumes into primary volumes and immediately suspends these volumes. Because the B volumes are the now the primary volumes and the A volumes are now the secondary volumes, the Global Copy state of the secondary volumeschanged from Target Copy Pending to Suspended.

To initiate failover processing of the B volumes to the A volumes, complete the following step.

Note: You can enter the commands either for a DS8000 model or for a DS6000 model. The storage image ID for the DS6000 model is different.

Enter the **failoverpprc** command to change the Global Copy state on the B volumes from secondary, target pending to primary, suspended.

dscli> failoverpprc -dev storage_image_ID -remotedev storage_image_ID SourceVolumeID:TargetVolumeID Example

```
dscli> failoverpprc -dev IBM.2107-75FA150 -remotedev IBM.2107-75FA120 0100:0100 0101:0101 0102:0102 0103:0103
```

Note: Unlike most other commands, the *storage_image_ID* and *SourceVolumeID:TargetVolumeID* values are reversed for the **failoverpprc** command because the command is issued directly to the remote server. You must place the values for the remote server first and the values for the original primary server second. This reversal (where the remote server is no longer the primary server) can have the following results:

- The source volumes become suspended
- The target volumes become suspended and they are available for mounting

An information message like the following one is generated for each Global Copy pair that is changed and moved to a suspended state.

Remote Mirror and Copy pair IBM.2107-75FA150/0100:IBM.2107-75FA120/0100 successfully suspended.

All B volumes must successfully process the **failoverpprc** command before you can proceed to the next step to analyze the FlashCopy relationships and enter the appropriate commands.

Analyzing and validating the consistency group state

Complete this task to analyze and validate the consistency group state. This is the fourth step in the Global Mirror failover recovery process.

Before you can initiate this task, you must ensure that the following tasks have been completed:

- Global Mirror processing has been ended at the primary server site.
- The status of Global Mirror transaction processing before the disaster caused the process to end has been obtained.
- Failover processing from B volumes to A volumes has completed with the B volumes state being changed from a secondary, target pending state to a primary, suspended state.

The consistency group state must be validated. This means that you must investigate whether all FlashCopy relationships are in a consistent state. Query the FlashCopy relationships that exist between B volumes and C volumes to determine the state of the FlashCopy relationship at the time that the primary server experienced a failure. Global Mirror might have been in the middle of forming a consistency group and FlashCopy might not have completed to create a complete set of consistent C volumes.

When you query a FlashCopy pair, there are two key pieces of information, that determine whether the C volume set is consistent or needs intervention to correct some states.

Revertible status

The revertible status is indicated as Enable or Disable and shows whether the FlashCopy is revertible or nonrevertible. A nonrevertible state means that a FlashCopy process has completed successfully and all changes are committed.

Sequence number

The sequence number indicates the number of the actual or last FlashCopy process (if the FlashCopy process is finished).

There are some combinations of revertible states and FlashCopy sequence numbers that require different corrective actions; this is, what you are looking for when you do your analysis.

Complete the following steps to analyze and validate the consistency group state. The example command displayed in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **lsflash** command to provide a report that lists the FlashCopy relationships and status information for each FlashCopy relationship in the list. Enter the **lsflash** command at the dscli command prompt with the following parameters and variables:

dscli> lsflash -dev storage_image_ID -l source_volume_ID:target_volume_ID

Example

dscli> lsflash -dev IBM.2107-75FA150 -l 0100:0200 0101:0201 0102:0202 0103:0203

Remember that your remote server has become your primary server.

2. Analyze your report to determine the state of the consistency group between the B volume and C volume. You are looking primarily at the FlashCopy relationships and your analysis determines your next step in the recovery process.

The following provides the types of FlashCopy relationships that might exist and a reference to the action that must be taken:

- The FlashCopy relationships are nonrevertible and all the sequence numbers are equal. Action: No further action is necessary and the set of C volumes is consistent and a complete copy.
- The FlashCopy relationships are revertible and all the sequence numbers are equal. Action: Issue the **revertflash** command to revert the FlashCopy pairs to the last consistency group.
- All the FlashCopy sequence numbers are equal and at least one of the FlashCopy relationships is nonrevertible. Action: Issue the **commitflash** command to commit the data to a target volume to form a consistency group between the source and target.
- The FlashCopy relationships appear as follows:
 - **a**. Some of the FlashCopy relationships completed processing so that a consistent group was created. These FlashCopy relationships are no longer revertible.
 - b. Some of the FlashCopy relationships have not completed creating a consistency group formation. They are still in a revertible state.
 - **c**. All the FlashCopy relationships have the same FlashCopy sequence number. This indicates that all the FlashCopy relationships were involved in the same consistency group.

Action: Issue the **commitflash** command to commit the data of the FlashCopy relationships that have not completed creating a new consistency group so that a consistency group is formed.

- There is a group of FlashCopy pairs that are all revertible and another group of FlashCopy pairs that are all nonrevertible. In addition, all the FlashCopy sequence numbers are not equal. However, the following conditions exists:
 - a. The FlashCopy sequence number for all revertible pairs is equal.
 - b. The FlashCopy sequence number for all nonrevertible pairs is equal.

Action: Issue the revertflash command to revert the FlashCopy pairs to the last consistency group.

• The FlashCopy sequence numbers are not equal for all FlashCopy relationships in the concerned consistency group and either *a* or *b* in the previous bullet was not true. This indicates that the consistency group is corrupted. **Action:** Contact your IBM service representative.

Note: When you know the state of all the FlashCopy relationships, you might want to initiate a tape backup of the C volume.

After determining the state of the FlashCopy relationships, issue the **revertflash** or **commitflash** commands, as appropriate.

Using the revertflash command to correct FlashCopy relationships

Complete this task to correct the revertible states and FlashCopy sequence numbers that require the use of the DS CLI **revertflash** command. This is the fifth step in the Global Mirror failover recovery process unless your corrections require the use of the **commitflash** DS CLI command. In this case, the use of the **commitflash** command becomes your fifth step.

You can use the **revertflash** command only when your analysis of the FlashCopy relationships reveals one of the following conditions:

- The FlashCopy relationships are revertible and all the sequence numbers are equal.
- There is a group of FlashCopy pairs that are all revertible and another group of FlashCopy pairs that are all nonrevertible. In addition, all the FlashCopy sequence numbers are not equal. However, the following conditions exist:
 - The FlashCopy sequence number for all revertible pairs is equal.
 - The FlashCopy sequence number for all nonrevertible pairs is equal.

The revert action removes the FlashCopy relationship changes and resets them to the last consistency group state. The revertible state is set to No.

Complete the following step to correct the applicable FlashCopy relationships. The example command in this task is shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

Issue the **revertflash** command to correct the FlashCopy relationships and reset them to the last consistency group state. Enter the **revertflash** command at the dscli command prompt with the following parameters and variables:

dscli> revertflash -dev storage_image_ID SourceVolumeID Example

dscli> revertflash -dev IBM.2107-75FA150 0100

Notes:

- 1. Remember that the *storage_image_ID* is the value for the remote server that has been designated the primary server until the original primary server is available for use.
- 2. Global Mirror operations have run the establish FlashCopy revertible processing as it was trying to form a consistency group before the disaster occurred. If your analysis, through use of the **lsflash** command, has determined that a **revertflash** command is needed, then there is no need to issue a new **mkflash** command.

A confirmation message like the following one is generated for each FlashCopy relationship that has been successfully reset.

FlashCopy pair 0100:0200 successfully reverted to the previous consistency.

After all the FlashCopy relationships have been corrected, you are ready to use the fast reverse restore process, which is the next step in the Global Mirror disaster recovery process.

Using the commitflash command to correct FlashCopy relationships

Complete this task to correct the revertible states and FlashCopy sequence numbers that require the use of the DS CLI **commitflash** command. This is the fifth step in the Global Mirror failover recovery process unless your corrections require the use of the **revertflash** DS CLI command. In this case, the use of the **revertflash** command becomes your fifth step.

You can use the **commitflash** command only when your analysis of the FlashCopy relationships reveals one of the following conditions:

- All the FlashCopy sequence numbers are equal and at least one of the FlashCopy relationships is nonrevertible.
- The FlashCopy relationships exist as follows:
 - Some of the FlashCopy relationships completed processing so that a consistent group has been created. These FlashCopy relationships are no longer revertible.
 - Some of the FlashCopy relationships have not completed creating a consistency group formation. These FlashCopy relationships are still in a revertible state.
 - All the FlashCopy relationships have the same FlashCopy sequence number. This indicates that all the FlashCopy relationships belong in the same consistency group.

The commit action keeps a FlashCopy relationship in its current state and resets the revertible state to No. When the **commitflash** command is processed, the data in these relationships is committed to the consistency group to which it would have become a part before the disaster occurred.

Complete the following step to correct the applicable FlashCopy relationships. The example command in this task is shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

Issue the **commitflash** command to correct the FlashCopy relationships and commit them to the consistency group that was being formed before the disaster occurred. Enter the **commitflash** command at the dscli command prompt with the following parameters and variables: dscli> commitflash -dev *storage_image_ID SourceVolumeID*

Example

dscli> commitflash -dev IBM.2107-75FA150 0100

Notes:

- 1. Remember that the *storage_image_ID* is the value for the remote server that has been designated the primary server until the original primary server is available for use.
- 2. Global Mirror operations have performed the establish FlashCopy revertible processing as the Global Mirror operations were trying to form a consistency group before the disaster occurred. If your analysis, through use of the **lsflash** command, has determined that a **commitflash** command is required, then there is no need to issue a new **mkflash** command.

A confirmation message like the following one is generated for each FlashCopy relationship that has been successfully reset.

FlashCopy pair 0100:0200 successfully committed.

After all the FlashCopy relationships have been corrected, you are ready to use the fast reverse restore process, which is the next step in the Global Mirror disaster recovery process.

Using fast reverse restore processing to create consistency

Complete this task to create the same consistent data on the B volumes that you have on the C volumes. The fast reverse restore option allows you to reverse a FlashCopy relationship without waiting for the background copy of a previous FlashCopy operation to finish. This is the sixth step in the Global Mirror disaster recovery process.

The fast reverse restore option reverses a FlashCopy target volume and allows consistent data to be copied back to its associated source volume, without having to wait for the background copy from the original source to the original target to complete. You can then vary the volumes online and start your applications.

Fast reverse restore processing creates a background copy of all tracks that changed on the B volume since the last consistency group formation. This results in the B volume becoming the same as the image that was present on the C volume. However, this process ends the FlashCopy relationship, so that the C volume is no longer usable.

Use the DS CLI **reverseflash** command with the **-fast** parameter to accomplish this task. This command results the following conditions:

- Start background copy from the C volumes to the B volumes.
- No change recording.
- There must be no I/O allowed to the B or C volumes during the fast reverse restore process.

Complete the following step to create the same consistent data on the B volumes that you have on the C volumes. The example command in this task is shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the command that is described in the steps below either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

Issue the **reverseflash** command to create the same consistency on the B volumes that you have on the C volumes. Enter the **reverseflash** command at the dscli command prompt with the following parameters and variables:

dscli> reverseflash -dev *storage_image_ID* -fast -tgtpprc *source_volume_ID:target_volume_ID* **Example**

dscli> reverseflash -dev IBM.2107-75FA150 -fast -tgtpprc 0200:0100

Notes:

- 1. The **-fast** parameter determines that the **reverseflash** command is processed before the background copy completes.
- 2. When you reverse a FlashCopy relationship, the source and target IDs that you specify for the **reverseflash** command must be the same as the source and target IDs from the original FlashCopy relationship (the source and target IDs from the relationship before it is reversed). The target volume ID is the value that is specified for the C volume. The data from this volume is copied back to the source volume ID, which is the B volume.
- **3**. The **-tgtpprc** parameter allows the FlashCopy target volume (B volume) to be a Remote Mirror and Copy source volume.
- 4. The *storage_image_ID* parameter is the value that is assigned to the remote storage unit, which has become the primary storage unit as a result of the failover action.

5. You must wait for the background copy to complete before you can proceed to the next process.

Waiting for the background copy to complete

Complete this task to determine when all fast reverse restore operations are complete and when no more FlashCopy relationships exist between the B volumes and the C volumes. This is the seventh step in the Global Mirror disaster recovery process.

The fast reverse restore operations complete the data transfer from the C volumes to the B volumes. However, before you can proceed with any additional steps, the back ground copy must complete. When the background copy is completed, FlashCopy relationships no longer exist between the B volumes and C volumes. Also, the C volume is no longer usable. Both of these operations must complete before you can move on in the disaster recovery process.

The best way to determine if these operations are complete is to periodically issue the **lsflash** command against the B volumes to query the existence of FlashCopy relationships.

Complete the following steps to determine that no FlashCopy relationships exist between the B volumes and the C volumes. The example command in this task is shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **lsflash** command to check the existence of FlashCopy relationships between the B volume and the C volume. Enter the **lsflash** command at the dscli command line prompt with the following parameters and variables:

```
dscli> lsflash -dev storage_image_ID -s target_volume_ID
```

```
Example
```

```
dscli> lsflash -dev IBM.2107-75FA150 -s 0200
```

Notes:

- a. The *storage_image_ID* is the manufacture, storage unit type, and serial number value of the remote storage unit that has become the primary unit because of the disaster.
- b. The **-s** parameter limits the report information that is returned only to the FlashCopy pair relationships that still exist.
- **c.** By designating only the target volume ID, you are further limiting the report to display just the target side of the FlashCopy pair relationship. When the report returns a blank screen, it indicates that background copy has completed and that no FlashCopy relationships exist between the B volume and the C volume.
- 2. Repeat Step one periodically until no FlashCopy relationships exist between the B volume and the C volume.

After the fast reverse restore and the background copy operations have completed, you can proceed to the next task which is reestablishing the FlashCopy relationship between the B volume and the C volume.

Reestablishing the FlashCopy relationships between B volumes and C volumes

Complete this task to reestablish the FlashCopy relationships between your B volumes and C volumes. This is the eighth step in using the Global Mirror disaster recovery process.

In this task, you are reestablishing the FlashCopy relationships between the B volumes and C volumes that were established for Global Mirror operations before the disaster occurred. The task is not much different than the one that you used to establish FlashCopy relationships during the set up of your Global Mirror environment.

Complete the following steps to create FlashCopy relationships between the B volumes and the C volumes. The example command in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **mkflash** command to create FlashCopy relationships at the remote site between the Global Copy secondary volumes and the FlashCopy target volumes. Enter the **mkflash** command at the dscli command prompt with the following parameters and variables:

```
dscli> mkflash -dev storage_image_ID -tgtinhibit -persist -record -nocp sourcevolumeID:targetvolumeID
```

Example

```
dscli> mkflash -dev IBM.2107-75FA150 -tgtinhibit -record -persist -nocp 0001:0004
```

Notes:

- a. Specify the secondary storage unit MTS (which has become the primary storage unit because of the disaster) for the *-dev storage_image_ID* parameter.
- b. Use the **-tgtinhibit** parameter to inhibit writes on the target volume.
- c. Use the -record parameter to activate change recording on the volume pair.
- d. Use the **-persist** parameter to retain the FlashCopy relationship after the background copy completes.
- e. Use the **-nocp** parameter to inhibit the background copy.
- f. The *source_volume_ID* is the value associated with the B volumes and the *target_volume_ID* is the value associated with the C volumes.
- **2**. Use the **lsflash** command to check the status of the FlashCopy relationships after you have processed the **mkflash** command.

After you have reestablished the FlashCopy relationships, you can start host I/O processing at the remote site on the B volumes. The production operation on the remote site, in this configuration, remains until you are ready to return production to the local site.

Preparing to reinstate production at the local site

Complete this task to begin the process of returning production to your local site. Just as there was a series of steps in the failover recovery process to your remote site, there are a series of steps that you must take to return production to your local site.

Returning production to its original implementation is called a failback recovery. After restoring operations at Site A, you can schedule a failback operation to synchronize data and to enable production to resume at your original site, Site A.

This task is initiated when your local site has been repaired and is operational. The first step in returning production to site A is to create Fibre Channel paths between Site B to Site A and between the specific LSSs.

Complete the following steps to create Fibre Channel paths from Site B to Site A and between the specific LSSs. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with values declared for the variables.

Notes:

- 1. Before you can establish the paths, you must obtain the worldwide node name that is associated with the remote storage unit. In this task your remote storage unit is your local storage unit (Site A).
- 2. You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.
- 1. Issue the **lssi** command against the Site A storage unit to obtain its worldwide node name. A report is displayed that provides the specific information about the Site A storage unit. Enter the **lssi** command at the dscli command prompt with the following parameters and variables:

dscli> lssi -l storage_image_ID

Example

dscli> lssi -l IBM.2107-75FA120

Record the worldwide node name because it is used in the next step.

2. Issue the **mkpprcpath** command to create the Fibre Channel paths from Site B to Site A and between the specific LSSs. Enter the **mkpprcpath** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprcpath -dev storage_image_ID -remotedev storage_image_ID -remotewwnn wwnn -srclss source_LSS_ID -tgtlss target_LSS_ID source_port_ID:target_port_ID

Example

```
dscli> mkpprcpath -dev IBM.2107-75FA150 -remotedev
IBM.2107-75FA120 -remotewwnn 12341234000A000A
-srclss IBM.2107-75FA150/00 -tgtlss IBM.2107-75FA120/01
I1A20:I2A10
```

Notes:

- a. The **-dev** parameter specifies the ID of your source storage unit. At this point in time, your source is the Site B storage unit.
- b. The **-remotedev** parameter specifies the ID of the secondary storage unit. At this point in time, the remote storage unit is your Site A storage unit.
- **c.** The **-remotewwnn** parameter must specify the worldwide node name of the secondary storage unit (Site A at this point in time). If you specify the worldwide node name of the primary storage unit (Site B), the command fails.
- d. The **-srclss** parameter refers to Site B storage unit as the source.
- e. The **-tgtlss** parameter specifies the Site A storage unit as the target.
- f. The *source_port_ID:target_port_ID* value has the Site B port ID as the source and the Site A port ID as the target.

After you have established the paths, you are ready to move on to the second step on the failback recovery process which involves issuing the **failbackpprc** command from the B volume to the A volume.

Resynchronizing the volumes

Complete this task to resynchronize the volumes, B volumes to A volumes. This is the second step in the failback recovery process that allows production to be returned to your A site.

Before you can do this task, you must ensure that you have created paths from Site B to Site A between the specific LSSs.

This task requires the use of the **failbackpprc** command. Processing this command resynchronizes the volumes in the following manner, depending on the volume state:

- If a volume at Site A is in simplex state, all of the data for that volume is sent from Site B to Site A.
- If a volume at Site A is in full-duplex or suspended state and without changed tracks, only the modified data on the volume at Site B is sent to the volume at Site A.
- If a volume at Site A is in a suspended state but has tracks that have been modified, the volume at Site B discovers which tracks were modified at any site and sends the following data from Site B to Site A:
 - The tracks that were changed on Site A.
 - The tracks that were marked at Site B.

Complete the following step to resynchronize your volumes. The example commands displayed in this task are shown in two formats. The first format shows the type of information the command requires. The second format provides the command with declared values for the variables.

Notes:

- 1. This task does not affect normal operations. All your operations remain at the remote site (Site B). This task is just part of the preparation you need to make to transfer operations back to Site A after it has been repaired.
- 2. You can issue the command that is described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

Issue the **failbackpprc** command to resynchronize your volumes. Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev storage_image_ID -remotedev storage_image_ID -type gcp source_volume_ID:target_volume_ID

Example

```
dscli> failbackpprc -dev IBM.2107-75FA150 -remotedev IBM.2107-75FA120
-type gcp 1000:1000
```

Notes:

- 1. The **-dev** parameter specifies the ID of your source storage unit. At this point in time, your source is the Site B storage unit.
- 2. The **-remotedev** parameter specifies the ID of the target storage unit. At this point in time, the remote storage unit is your Site A storage unit.
- **3**. The *source_volume_ID:target_volume_ID* value has the Site B volume ID as the source and the Site A volume ID as the target.

After submitting this command for processing, you must track the progress of the transaction until it completes its first pass. So, querying for first pass completion is the next step in the failback recovery process.

Querying, quiescing, and re-querying

Complete this task series to query for the first pass of the out-of-sync bitmap completion, to quiesce your system, and to complete the query process to ensure that the out-of-sync tracks equal 0. This series of tasks is the third step in the failback recovery process that allows production to be returned to your A site.

To complete this series of tasks, you must ensure that you have resynchronized the volumes, B volumes to A volumes.

This series of tasks requires the use of the lspprc command and that you quiesce your system.

Complete the following steps to complete the third step of the failback recovery process. The example commands in this task are shown in two formats. The first format shows the type of information that the command requires. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **1spprc** command periodically to identify when the first pass of the out-of-sync (OOS) bitmap completes. Depending on the number of transactions that you have, some time elapses. Enter the **1spprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev *storage_image_ID* -remotedev *storage_image_ID* -l *SourceVolumeID:TargetVolumeID* Example

dscli> lspprc -dev IBM.2107-75FA150 -remotedev IBM.2107-75FA120 1000:1000

Notes:

- a. The **-dev** parameter specifies the ID of your source storage unit. Your source is the Site B storage unit.
- b. The **-remotedev** parameter specifies the ID of the target storage unit. The remote storage unit is your Site A storage unit.
- c. The *source_volume_ID:target_volume_ID* value has the Site B volume ID as the source and the Site A volume ID as the target.
- 2. Quiesce your I/O and unmount your file systems at the B site to preserve the integrity of your file system.

Note: Unmounting your file systems flushes the host cache and ensures that you actually copy valid data sets.

3. Reissue the **"lspprc" on page 455** command periodically to identify when the remaining bits completely drain from the B site. This is indicated when the out-of-sync (OOS) tracks equal zero. Depending on the number of transactions that you have, some time elapses. Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID -l SourceVolumeID:TargetVolumeID

Example

dscli> lspprc -dev IBM.2107-75FA150 -remotedev IBM.2107-75FA120 1000:1000

After this task is completed, you are ready to establish the remote mirror and copy paths from Site A to Site B.

Reestablishing remote mirror and copy paths (site A to site B)

Complete this task to reestablish the remote mirror and copy paths between site A and site B in preparation for the transfer of operations from the B site to the A site. This is the fourth step in the failback disaster recovery process.

Each of the prior tasks must be competed in sequence in order for this task to succeed.

This task is like your initial creation of remote mirror and copy paths while setting up your Global Mirror environment, before the disaster occurred.

Create paths so that the logical subsystems (LSSs) are associated with each other. These are the ports that the copy services I/O pass through. It is preferred that they are not the same ports that are used for host I/O. This ensures that there is enough capacity for the data transfer.

Complete the following steps to create remote mirror and copy paths between all Global Mirror source and target pairs and between the master and subordinate storage units. The example commands in this task are shown in two formats. The first format shows the type of information that is required by the command. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

 Obtain the worldwide node name of the secondary storage unit. This information is needed when you do the next step. Enter the lssi or showsi at the dscli command prompt as follows: dscli> lssi -1

This is the entire command. No additional variables are needed.

The **showsi** command does contain a variable and a command: dscli> showsi *storage_image_id* -fullid

Example

dscli> showsi -fullid IBM.2107-75FA120

Notes:

- a. Use the **lssi** command if you want to display a list of all the storage image instances for a storage complex and the status information for each storage image in the list.
- b. Use the showsi command if you want to display the detailed properties of a specific storage unit.
- **c**. Use the **-fullid** DS CLI command with the **showsi** command to display fully qualified IDs, which include the storage image ID, for every ID that is displayed in the command output.
- d. Record the worldwide node name for the secondary (target) site B storage unit so that you can use it when you issue the **mkpprcpath** command.
- 2. Issue the **mkpprcpath** command to create the remote mirror and copy paths between all Global Mirror source and target pairs and between the master and subordinate storage units. Enter the **mkpprcpath** command at the dscli command prompt with the following parameters and variables as follows:

dscli> mkpprcpath -dev storage_image_ID -remotedev storage_image_ID -remotewwnn wwnn -srclss source_LSS_ID -tgtlss target_LSS_ID source_port_ID:target_port_ID

Example

```
dscli> mkpprcpath -dev IBM.2107-75FA120 -remotedev
IBM.2107-75FA150 -remotewwnn 12341234000A000F
-srclss 00 -tgtlss 01 I1A10:I2A20
```

Notes:

a. The -remotedev parameter specifies the ID of the secondary storage unit.

- b. The **-remotewwnn** parameter must specify the worldwide node name of the secondary storage unit. If you make a mistake and specify the worldwide node name of the primary storage unit, the command fails.
- c. The shortened version of the -srclss parameter is shown (value = 00) because the example uses the fully qualified -dev storage_image_ID parameter. If the fully qualified -dev parameter is not used, you must use the fully qualified -srclss source_LSS_ID value. For example: -srclss IBM.2107-75FA120/00.
- d. The shortened version of the -tgtlss parameter is shown (value = 01) because the example uses the fully qualified -dev storage_image_ID parameter. If the fully qualified -dev parameter is not used, you must use the fully qualified -tgtlss target_LSS_ID value. For example: -tgtlss IBM.2107-75FA120/01.

e. The shortened version of the *source_port_ID:target_port_ID* parameter is shown (value = I1A10:I2A20) because the example uses the fully qualified -dev *storage_image_ID* and -remotedev *storage_image_ID* parameters. If the fully qualified -dev and -remotedev parameters are not used, you must use the fully qualified source_port_ID:target_port_ID value. For example: IBM.2107-75FA120/I1A10:IBM.2107-75FA150/I2A20.

The fully qualified *source_port_ID:target_port_ID* parameter is positional on your command line. It must be placed after the **-tgtlss** parameter and value. For example:

dscli> mkpprcpath -srclss 00 -tgtlss 01 IBM.2107-75FA120/I1A10:IBM.2107-75FA150/I2A20

Running Global Copy failover processing to the A volumes

Complete this task to run failover processing from A volumes to B volumes so that the A volumes become the primary volumes and B volumes become the secondary volumes. This is the fifth step in the recovery failback process.

The resynchronization of the A volumes and B volumes must be completed (no out-of-sync tracks) before you can proceed with this task.

You must issue this restore failover request on the Global Copy volumes pair to reestablish the extended distance relationship and create the A Volume to B Volume Global Copy relationship.

Complete the following step to complete failover processing from the A volumes to the B volumes. The example commands in this task are shown in two formats. The first format shows the type of information that is required by the command. The second format provides the command with declared values for the variables.

Note: You can issue the command that is described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

Issue the **failoverpprc** command to reestablish the extended distance relationship and create the A Volume to B Volume Global Copy relationship. Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failoverpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp SourceVolumeID:TargetVolumeID

Example

dscli> failoverpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -type gcp 0100:0100 0101:0101 0102:0102 0103:0103

Note: The SourceVolume_ID is the A volume and the TargetVolume_ID is the B volume. A confirmation message like the following example is generated for each Global Copy pair that has been changed and moved to a state of suspended.

PPRC pair IBM.2107-75FA120/0100:IBM.2107-75FA150/0100 successfully suspended.

Note: All A volumes must successfully process the **failoverpprc**command before you can proceed to the next step.

Running Global Copy failback processing for the A volumes

Complete this task to run Global Copy failback processing for the A volumes. This process resynchronizes the volumes at Site A with volumes at Site B and restarts mirroring from site A (local site) to site B (remote site). This is the sixth step in the recovery failback process.

The failover processing that is described in the fifth step of the recovery failback process must have completed. The failover process converted the full-duplex target volumes at site A to suspended source volumes. The volumes at site A started the change recording process while in failover mode.

The failback processing that is described in this task can be issued 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.

Complete the following step to complete failback processing for the A volumes. The example commands in this task are shown in two formats. The first format shows the type of information that is required by the command. The second format provides the command with declared values for the variables.

Note: You can issue the command that is described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

Issue the **failbackpprc** command to resynchronize your volumes. Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev storage_image_ID -remotedev storage_image_ID -type gcp SourceVolume_ID:TargetVolume_ID

Example

```
dscli> failbackpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -type gcp 1000:1000
```

Notes:

- 1. The -dev parameter specifies the ID of your source storage unit. Your source is the Site A storage unit.
- 2. The **-remotedev** parameter specifies the ID of the target storage unit. The remote storage unit is your Site B storage unit.
- **3**. The *SourceVolume_ID:TargetVolume_ID* value has the Site A volume ID as the source and the Site B volume ID as the target.

Resuming Global Mirror processing at site A

Complete this task to start or resume Global Mirror processing at the A site. This is the final step in the failback recovery process.

Site A has been repaired and connectivity reestablished with your remote site. You have followed all the prior steps in sequence and are ready to start up I/O processing from your local site.

Complete the following steps to start or resume Global Mirror processing. The example commands in this task are shown in two formats. The first format shows the type of information that is required by the command. The second format provides the command with declared values for the variables.

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Issue the **mkgmir** command to start Global Mirror processing. Enter the **mkgmir** command at the dscli command prompt using the following parameters and variables:

```
mkgmir -dev storage_image_ID -lss LSS_ID -cginterval seconds
-coordinate milliseconds -drain seconds -session session_ID
Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID
```

Example

```
dscli> mkgmir -dev IBM.2107-75FA120 -lss 10 -cginterval 0 -coordinate 50 -drain 30 -session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00
```

Note: Issuing the **mkgmir** command requires that you specify the tuning parameters. The values for the tuning parameters are not retained when you end Global Mirror processing. So, in the case where you need to change the Global Mirror topology parameters, you need to resubmit the tuning parameters when you restart Global Mirror processing.

2. Issue the **resumegmir** command to continue Global Mirror processing after you have paused Global Mirror processing. Enter the **resumegmir** command at the dscli command prompt using the following parameters and variables:

dscli> resumegmir -dev storage_image_ID -lss LSS_ID -session session_ID Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID

Example

dscli> resumegmir -dev IBM.2107-75FA120 -lss 10 -session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00

Note: You might want to change or maintain the values that you had on your B site for the tuning parameters. You must restate these values before you process the **resumegmir** command. You cannot state a value for just one of the tuning parameters. You must restate all of the values (**-cginterval**, **-coordinate**, and **-drain**). The following example shows how to enter the **resumegmir** command to provide these values:

dscli> resumegmir -dev IBM.2107-75FA120 -lss 10 -cginterval 5 -coordinate 50 -drain 30 -session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00

Chapter 9. Recovery scenarios for planned and unplanned outages using Metro/Global Mirror

Information is provided to help use Metro/Global Mirror functions during planned and unplanned outages.

In a Metro/Global Mirror environment (DS8000 only) that uses failover and failback operations, you can switch applications to an alternate site during planned and unplanned outages. Planned outages are events such as maintenance updates. Unplanned outages are events such as disasters.

Several scenarios, including the following, are described:

- Setting up a Metro/Global Mirror environment
- Failover and restore operations to the intermediate (B) site
- Failover and restore operations to the remote (C) site

Note: The steps in these scenarios are examples. Other configurations might be possible but might not be supported by IBM.

Setting up a Metro/Global Mirror environment

Use this process to set up your system environment (DS8000 only) to use Metro/Global Mirror.

The volumes at the local and intermediate sites must be connected to the site from which the commands are going to be issued. For example, if the intermediate site volumes are connected to an intermediate site storage unit only, the Global Mirror setup commands are issued there. If the local site has connectivity to the intermediate site and local site volumes, then you can issue the commands from the local site.

Configure the following Metro/Global Mirror environment, which uses three sites (local, intermediate, and remote) and a minimum of four volumes (volume A, volume B, volume C, and volume D) on three storage units. For ease of description, the Metro/Global Mirror configuration is described in terms of A, B, C, and D volumes. Some environments can contain hundreds or thousands of volumes.

Complete the following steps in sequence. You can issue the commands in any order with the following two exceptions:

- You must establish paths before you can establish pairs or start a session.
- You must create a session to an LSS before you can add a volumes to the session.

In the event of a disaster, external automated software can detect the loss of the local site. (Automation software is not provided with the storage unit; it must be supplied by the user.) Data consistency on the target volumes must be managed using either automation procedures that are able to freeze the activity on the required volumes in the event of an unexpected outage or by remote mirror and copy operations that use freeze and run commands.

Complete the following steps to set up your Metro/Global Mirror environment. The command parameters and output are provided as examples.

- 1. Ensure that the storage units are configured, assigned, and operating in a normal state.
- 2. Identify all volumes that will participate in a session. Identify which volumes are to be source and target volumes for Metro Mirror, Global Copy, and FlashCopy relationships, and the storage unit that you will designate as the master storage unit.
- 3. At each site, establish Fibre Channel paths.

- a. Determine that there are I/O ports available for paths between the source and the target LSSs using the lsavailpprcport command. See "Determining which I/O ports are available for paths" on page 579 for more information.
- b. Set up paths between local and intermediate sites for the Metro Mirror volume pairs. Enter the mkpprcpath command at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprcpath -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P
-remotewwnn 5005076303FFC550 -srclss 61 -tgtlss 63 -consistgrp
I0102:I0031 I0002:I0102
```

The following represents an example of the output:

CMUC00149I mkpprcpath: Remote Mirror and Copy path 61:63 successfully established.

Ensure that you create the paths with the PPRC consistency group option (using the **-consistgrp** parameter) for the A and B volume pairs in Metro Mirror relationships. Specifying the consistency group option ensures that volume pairs from the specified LSSs that share the same paths belong to this consistency group. When an error occurs that affects any of these volumes in the consistency group, the volumes in the consistency group become suspended and enter a long-busy state until a consistency group operation is run. See "Defining a path that has the consistency option enabled" on page 576 for more information.

c. Set up paths between the intermediate and remote sites for the Global Copy volume pairs. Enter the **mkpprcpath** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprcpath -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 -remotewwnn 5005076303FFC220 -srclss 62 -tgtlss 64 I0033:I0303

The following represents an example of the output: CMUC00149I mkpprcpath: Remote Mirror and Copy path 64:62 successfully established.

See "Creating remote mirror and copy paths" on page 567 for more information.

- d. Use the **lspprcpath** command to view the newly created paths. See "Displaying the status of established paths" on page 566 for more information.
- 4. At the intermediate site, create Global Copy volume pairs between the intermediate and remote sites. Create the pairs from the intermediate storage unit to the remote storage unit using the previously established paths. Ensure that you specify the **-cascade** parameter to allow the source volume in a Global Copy relationship to be eligible to be a target volume for another relationship at the same time.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 -type gcp -mode nocp -cascade 0700-075f:1200-125f

The following represents an example of the output:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0700:1200 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0701:1201 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0702:1202 successfully created.

See "Creating a Global Copy relationship" on page 572 for more information.

- 5. At the local site:
 - a. Establish Metro Mirror volume pairs between the local and intermediate sites, with the Incremental Resynchronization option enabled. Create the pairs from the local storage unit to the intermediate storage unit using the previously established paths. Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type mmir -mode full -incrementalresync enable 0700-075f:1200-125f

The following represents an example of the output:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0700:1200
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0701:1201
successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0702:1202
successfully created.

See "Creating a Metro Mirror relationship" on page 569 for more information.

6. At the remote site, create FlashCopy relationships. Create the relationships at the remote site between volume C (the FlashCopy source volume that is also the target volume of the Global Mirror session) and volume D (the FlashCopy target volume).

Note: Delay initial FlashCopy operations until the Global Copy pairs have completed their first pass of the copying process.

Enter the **mkflash** command at the dscli command prompt with the following parameters and variables:

```
dscli> mkflash -dev IBM.2107-75ALAG1 -tgtinhibit -record -persist -nocp 1200-125f:1900-195f
```

The following represents an example of the output:

CMUC00137I mkflash: FlashCopy pair 1200:1900 successfully created. CMUC00137I mkflash: FlashCopy pair 1201:1901 successfully created. CMUC00137I mkflash: FlashCopy pair 1202:1902 successfully created. CMUC00137I mkflash: FlashCopy pair 1203:1903 successfully created.

When you create FlashCopy relationships, select the following options:

Enable Change Recording

Select this option to activate the change recording feature on the volume pair that is participating in a FlashCopy relationship.

Note: The Persistent FlashCopy option is automatically selected because it is required with the Enable Change Recording option.

Inhibit writes to target volume

Select this option to ensure that updates cannot be made to the target volume. This ensures data consistency on the target volume. If you select the Inhibit writes to target option, the change recording feature is not active on the target volume.

Attention: Do not select the **Initiate background copy** option. This ensures that data is copied from the source volume to the target volume only if a track on the source volume is modified. See "Creating FlashCopy relationships (Global Mirror setup)" on page 594 for more information.

7. At the intermediate site:

a. **Define the Global Mirror session.** Define the same session on the LSS that contains the master and on every LSS that contains volumes to be added to the Global Mirror session.

Enter the **mksession** command at the dscli command prompt with the following parameters and variables:

dscli> mksession -lss 07 1

The following represents an example of the output:

CMUC00145I mksession: Session 1 opened successfully

See "Creating the Global Mirror session" on page 595 for more information.

b. Add volumes to the Global Mirror session. Enter the chsession command at the dscli command prompt with the following parameters and variables:

dscli> chsession -lss 07 -action add -volume 0700-075f 1 The following represents an example of the output: CMUC00147I chsession: Session 1 successfully modified.

See "Adding volumes to a session (Global Mirror)" on page 584 for more information.

Note: If you have many volumes that you want to add to a new or existing Global Mirror session, you might consider adding them to the session in stages. When you add a large number of volumes at once to an existing Global Mirror session, the available resources for Global Copy processing within the affected ranks might be used by the initial copy pass of the new volumes. New volumes on the same ranks as existing volumes can use all the processing resources for the initialization of the new volumes.

To avoid too much impact on your processing, you might consider adding new volumes to a Global Mirror session in small numbers per rank and wait until the first pass has completed before adding more volumes.

c. Start the Global Mirror session. The master storage unit begins forming consistency groups for the specified Global Mirror session.

Enter the **mkgmir** command at the dscli command prompt with the following parameters and variables:

dscli> mkgmir -lss 07 -session 1

The following represents an example of the output:

CMUC00162I mkgmir: Global Mirror for session 1 successfully started.

See "Starting Global Mirror processing" on page 590 for more information.

d. Issue a query to confirm that the session exists. Confirm that the individual LSS sessions are populated with the appropriate volumes. See "Querying Global Mirror processing" on page 588 for more information.

Enter the **showgmir** command at the dscli command prompt with the following parameters and variables:

dscli> showgmir 07

The following represents an example of the output (in table format for clarity):

ID	Master Count	Master Session ID	Copy State	Fatal Reason	CG Interval Time (sec- onds)	Coord. Time (milli- sec- onds)	Max CG Drain Time (sec- onds)	Current Time
IBM. 2107- 75ALA 2P/07	1	0x31	Running	Not Fatal	0	50	30	02/19/ 2006 10:34:51 MST

CG Time	Succes- sful CG Percen- tage	Flash- Copy Sequ- ence Number	Master ID	Subor- dinate Count	Master /Subor- dinate Assoc
02/19/ 2006 10:34:51 MST	50	0x4357D -504	IBM.2107 -75ALA 2P	0	-

Failover and restore operations to the intermediate site during a planned outage

Use this process to run failover and restore operations (DS8000 only) to the intermediate (B) site during an unplanned outage.

For this scenario, assume that you have to shut down the local site for any reason and move production from the local site to the intermediate site and then return production back to the local site. You can use the following failover and failback procedures for this scenario. It is assumed that you established Global Mirror sessions that are creating consistency groups at the local site and sending them to the remote site.

During the outage and until you resume processing at the local site, you run a failover operation to allow operations to run from your intermediate site, which is protected by a two-site Global Mirror configuration. Global Mirror continues sending updates to the storage unit at the remote site and continues to form consistency groups. When production is ready to return to the local site, you run a failback operation.

Note: When a local site fails, systems must be reset and subsequently restarted using data from the B volumes following a failover operation. GDPS HyperSwap can do this transparently (without any system outage for systems running at the intermediate site) through the use of a *single* script statement for planned outages and autonomically for unplanned outages.

Complete these tasks for failover and restore operations at the intermediate site: (The steps in this scenario are examples.)

- 1. At the local site, ensure that data consistency is achieved between the A to B volume pairs. This process helps coordinate the A volumes and B volumes consistency and allows consistent data to be copied to the remote site. You can use either of the following methods to create data consistency:
 - Quiesce I/O processing to the A volumes at the local site. Continue to step 2 on page 626.
 - Freeze write activity to the Metro Mirror primary volumes by completing the following steps:
 - a. **Freeze updates to the A volumes in Metro Mirror relationships across the affected LSSs.** Enter the **freezepprc** command at the dscli command prompt with the following parameters and variables:

dscli> freezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07-12

The following represents an example of the output:

CMUC00161W freezepprc: Remote Mirror and Copy consistency group 07:12 successfully created.

This process ensures that the B volumes are consistent at the time of the freeze. (One command per storage unit or LSS is required.) As a result of the freeze action, the following actions are taken:

- The established paths between the logical subsystem (LSS) pairs are deleted.
- The volume pairs that are associated with the source and target LSSs are suspended. During this time, the storage unit collects data that is sent to the A Metro Mirror volumes.
- I/O processing to the Metro Mirror volume pairs is temporarily queued during the time that updates are frozen.
- **b.** If wanted, you can view the state of the pair status at the local site after the **freezepprc** command has been processed. Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -fmt default 0700-075f

The following represents an example of the output:

Notes:

- 1) The command example uses the command parameter **-fmt default**. This command parameter specifies that the output be set to a space-separated plain text table.
- 2) The following table format is presented for clarity. The actual report is not displayed in this format.
- **3**) The report example represents the information that is reported on when you do not specify the **-1** parameter.

ID	State	Reason	Туре	Source- LSS	Timeout (secs)	Critical Mode	First Pass Status
0700:1200	Suspend- ed	Freeze	Metro Mirror	07	unknown	Disabled	Invalid
0701:1201	Suspend- ed	Freeze	Metro Mirror	07	unknown	Disabled	Invalid
0702:1202	Suspend- ed	Freeze	Metro Mirror	07	unknown	Disabled	Invalid

See "Viewing information about Metro Mirror relationships" on page 578 for more information.

c. Resume operations following a freeze.

Issue the **unfreezepprc** command to allow I/O processing to resume for the specified volume pairs. Enter the **unfreezepprc** command at the dscli command prompt with the following parameters and variables:

Note: This activity is sometimes referred to as a *thaw* operation. dscli> unfreezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07:12

The following represents an example of the output: CMUC00198I unfreezepprc: Remote Mirror and Copy pair 07:12 successfully thawed.

2. Issue a failover command to the B to A volume pairs. This process detects that the B volumes are cascaded volumes at the intermediate site. When the command processes, the B volumes remain as primaries in a duplex pending state and secondaries to the A volumes. The B volumes remain nonexistent (or unavailable) secondary volumes to the A volumes in a Metro Mirror relationship. (In a cascaded relationship, the B volumes cannot be primary volumes in a Metro Mirror and Global Copy relationship at the same time.) When the direction of the volumes are switched and I/O processing is directed to the *new* primary B volumes, it is essential that the primary volumes (the A volumes) be the same size as the secondary volumes (the B volumes).

See "Running a recovery failover operation" on page 578 for more information.

Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failoverpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -type gcp -cascade 1200-125f:1a00-1a5f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1200:1A00 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1201:1A01 successfully reversed.

- **3**. **Redirect host I/O processing to the B volumes.** Changes are recorded on the B volumes until the A volumes can be resynchronized with the B volumes.
- 4. When the A volumes are ready to return to production, pause the Global Mirror session between the B to C volumes. Direct this command to the same LSS that you used to start the session. This step is needed to later change the direction of the B volumes and restore the A volumes. Enter the **pausegmir** command at the dscli command prompt with the following parameters and variables:

dscli> pausegmir -dev IBM.2107-75ALA2P -quiet -lss 07 -session 1

The following represents an example of the output:

CMUC00165I pausegmir: Global Mirror for session 1 successfully paused.

See "Pausing Global Mirror processing" on page 589 for more information.

5. **Suspend (pause) the B to C volume pairs.** Because the site B volumes cannot be source volumes for Metro Mirror and Global Copy relationships, you must suspend the B to C volumes so that B to A volumes can be established. This step stops all incoming write I/O operations to the affected B and C volume pairs and helps prepare for a later resynchronization of the A volumes with the current operating B volumes.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 1200-125f:0700-075f

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

See "Pausing a Metro Mirror relationship" on page 572 for more information.

6. Establish paths between the local site LSS and intermediate site LSS that contain the B to A Metro Mirror volumes. Enter the mkpprcpath command at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprcpath -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P
-remotewwnn 5005076303FFC550 -srclss 07 -tgtlss 12 -consistgrp
I0102:I0031 I0002:I0102
```

The following represents an example of the output:

CMUC00149I mkpprcpath: Remote Mirror and Copy path 07:12 successfully established.

See "Creating remote mirror and copy paths" on page 567 for more information.

7. Issue a failback command to the B volumes (with A volumes as secondaries). Host I/O processing continues uninterrupted to the B volumes as the A volumes are made current. This command copies the changes back to the A volumes that were made to the B volumes while hosts are running on the B volumes. (In a DS CLI environment, where the local and intermediate sites use different management consoles, you have to use a different DS CLI session for the management console of the B volumes at the intermediate site.) See "Running a recovery failback operation" on page 577 for more information.

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -type gcp 1200-125f:1a00-1a5f

The following represents an example of the output:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1200:1a00 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1201:1a01 successfully failed back.

8. Wait for the copy process of the B to A volumes to reach full duplex status (all out-of-sync tracks have completed copying). Host writes are no longer tracked. You can monitor when the number of out-of-sync tracks reaches zero by querying the status of the volumes. See "Viewing information about Metro Mirror relationships" on page 578 for more information.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -l 1200-125f:1a00-1a5f

The following represents an example of the output:

ID	State	Reason	Туре	Out Of Sync Tracks	Tgt Read	Src Cascade
1200:1a00	Copy Pending	-	Metro Mirror	46725	Disabled	Disabled
1201:1a01	Copy Pending	-	Metro Mirror	46725	Disabled	Disabled

Tgt Cascade	Date Sus pend ed	Source LSS	Timeout (secs)	Crit Mode	First Pass Status	Incre- mental Resync	Tgt Write
Invalid	-	10	Unknown	Disabled	Invalid	Enabled	Enabled
Invalid	-	10	Unknown	Disabled	Invalid	Enabled	Enabled

9. Quiesce host I/O processing to the B volumes.

10. Issue a failover command to the A to B volume pairs. This process ends the B to A volume relationships and establishes the A to B volume relationships. Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> failoverpprc -dev IBM.2107-130126X -remotedev IBM.2107-75ALA2P
-type mmir 1a00-1a5f:1200-125f
```

The following represents an example of the output:

```
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200
successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201
successfully reversed.
```

See "Running a recovery failover operation" on page 578 for more information.

11. After the failover operation, you can view the status of the volumes with the lspprc command. Enter the lspprc command at the dscli command prompt with the following parameters and variables:

```
dscli> lspprc -dev IBM.2107-130126X -remotedev IBM.2107-75ALA2P -fmt default 1a00-1a5f
```

The following represents an example of the output:

Notes:

- **a**. The command example uses the command parameter **-fmt default**. This command parameter specifies that the output be set to a space-separated plain text table.
- b. The following table format is presented for clarity. The actual report is not displayed in this format.
- **c.** The report example represents the information that is reported on when you do not specify the **-1** parameter.

ID	State	Reason	Туре	Source- LSS	Timeout (secs)	Critical Mode	First Pass Status
0700:1200	Suspend- ed	Host Source	Metro Mirror	1A	unknown	Disabled	Invalid
0701:1201	Suspend- ed	Host Source	Metro Mirror	1A	unknown	Disabled	Invalid

ID	State	Reason	Туре	Source- LSS	Timeout (secs)	Critical Mode	First Pass Status
0702:1202	Suspend- ed	Host Source	Metro Mirror	1A	unknown	Disabled	Invalid

12. Reestablish paths (that were disabled by the freeze operation) between the local site LSS and intermediate site LSS that contain the B to A Metro Mirror volume pairs. Enter the mkpprcpath command at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprcpath -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P
-remotewwnn 5005076303FFC550 -srclss 07 -tgtlss 12
-consistgrp I0102:I0031 I0002:I0102
```

The following represents an example of the output:

CMUC00149I mkpprcpath: Remote Mirror and Copy path 07:12 successfully established.

See "Reestablishing remote mirror and copy paths (site A to site B)" on page 615 for more information.

13. Issue a failback command to the A to B volumes. This failback command completes the restoration of the A to B volume relationships (the B volume becomes the target). The replication of the data starts immediately when the command is finished. Depending on how many tracks have changed during the disaster recovery test, resynchronization might take a long time.

Note: At this point, you can resume host I/O processing to the local site if optimizing host availability is critical. However, new host I/O that is written to the A volumes at the local site is not fully protected by Global Mirror processing until the Global Mirror operation is restored in step 16 on page 630.

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> failbackpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type
mmir la00-la5f:1200-125f
```

The following represents an example of the output:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A00:1200 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A01:1201 successfully failed back.

14. **Reestablish Global Copy relationships between the B to C volumes with the -cascade option.** When the failback operation has been done, Global Copy relationships can be re-created.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 -type gcp -mode nocp -cascade 1200-125f:0700-075f

The following represents an example of the output:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1200:0700 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1201:0701 successfully created.

15. Wait until the first pass of the Global Copy copying processing of the B to C volume pairs has completed. You can monitor this activity by querying the status of the volumes.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 -l -fmt default 1200-125f:0700-075f

The following represents an example of the output:

ID	State	Reason	Туре	Out Of Sync Tracks	Tgt Read	Src Cascade
0700:1200	Copy Pending	-	Global Copy	0	Disabled	Disabled
0701:1201	Copy Pending	-	Global Copy	0	Disabled	Disabled
0702:1202	Copy Pending	-	Global Copy	0	Disabled	Disabled

Tgt Cascade	Date Sus pended	Source LSS	Timeout (secs)	Crit Mode	First Pass Status	Incre mental Resync	Tgt Write
Invalid	-	07	Unknown	Disabled	True	Enabled	Enabled
Invalid	-	07	Unknown	Disabled	True	Enabled	Enabled
Invalid	-	07	Unknown	Disabled	True	Enabled	Enabled

16. Resume Global Mirror. Now that the original infrastructure has been restored, you can resume the Global Mirror session.

Enter the **resumegmir** command at the dscli command prompt with the following parameters and variables:

dscli> resumegmir -dev IBM.2107-75ALA2P -session 1 -lss 07

The following represents an example of the output:

CMUC00164I resumegmir: Global Mirror for session 1 successfully resumed.

See "Resuming Global Mirror processing" on page 589 for more information.

17. **Resume host I/O processing to the A volumes.** Direct host I/O processing back to the A volumes in preparation for resuming host I/O on the A volumes.

18. Verify that consistency group are forming successfully.

Enter the **showgmir** -**metrics** command at the dscli command prompt with the following parameters and variables:

dscli> showgmir -metrics 07

The following represents an example of the output:

See "Querying Global Mirror processing" on page 588 for more information.

ID	Total Failed CG Count	Total Succes- sful CG Count	Succes- sful CG Percen- tage	Failed CG after Last Success	Last Succes- sful CG Form Time	Coord. Time (milli- seconds)	CG Interval Time (sec- onds)
IBM.2107 -130165X /07	0	55	100	0	02/20/ 2006 11:38:25 MST	50	0

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
30	-	-	No Error	-	-	-	-

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
No Error	-	-	-	-	No Error	-	-

Failover and restore operations to the intermediate site during an unplanned outage

Use this process to run failover and restore operations (DS8000 only) to the intermediate (B) site during an unplanned outage.

This scenario provides a disaster recovery solution if a failure occurs at your local site. You can run your operations from your intermediate site, which is protected by a two-site Global Mirror configuration, until your local site has recovered. Global Mirror continues sending updates to the storage unit at the remote site and continues to form consistency groups.

Complete these subtasks for failover and surviving restore operations at the intermediate site: (The steps in this scenario are examples.)

1. At the local site, ensure that data consistency is achieved between the A and B volumes. If the local site was not completely destroyed, it is essential that data from any surviving A and B volume pairs be copied and a consistent copy be achieved at the remote site. You can use freeze and unfreeze commands that are supported using external automation software to create data consistency to multiple Metro Mirror volume pairs.

To freeze write activity to Metro Mirror primary volumes, complete the following steps:

a. Freeze updates to the A volumes in Metro Mirror relationships across the affected LSSs. This ensures that the B volumes are consistent at the time of the freeze process. (One command per LSS is required.) Enter the **freezepprc** command at the dscli command prompt with the following parameters and variables:

dscli> freezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07-12

The following represents an example of the output:

CMUC00161W freezepprc: Remote Mirror and Copy consistency group 07:12 successfully created.

As a result of the freeze action, the following actions processing occurs:

- The established paths between the LSS pairs are deleted.
- The volume pairs that are associated with the source and target LSSs are suspended. During this time, the storage unit collects data that is sent to the A volumes in Metro Mirror relationships.
- I/O to the Metro Mirror volume pairs is temporarily queued during the time that updates are frozen.
- b. Resume operations following a freeze.

Issue the **unfreezepprc** command to allow I/O activity to resume for the specified volume pairs. Enter the **unfreezepprc** command at the dscli command prompt with the following parameters and variables:

Note: This activity is sometimes referred to as a *thaw* operation. dscli> unfreezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07:12 The following represents an example of the output:

CMUC00198I unfreezepprc: Remote Mirror and Copy pair 07:12 successfully thawed.

2. Issue a failover command to the B to A volumes. This process detects that the B volumes are cascaded volumes at the intermediate site. Enter the failoverpprc command at the dscli command prompt with the following parameters and variables:

```
dscli> failoverpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -type
gcp -cascade 1200-125f:1a00-1a5f
```

The following represents an example of the output:

```
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1200:1A00
successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1201:1A01
successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1202:1A02
successfully reversed.
```

See "Running a recovery failover operation" on page 578 for more information.

When the direction of the volumes are switched and I/O processing is directed to the *new* primary B volumes, it is essential that the primary volumes (the A volumes) be the same size as the secondary volumes (the B volumes).

- **3**. **Redirect host I/O processing to the B volumes.** Changes are recorded on the B volumes until the A volumes can be resynchronized with the B volumes.
- 4. When the A volumes are ready to return, pause the Global Mirror session between the B to C volumes. Direct this command to the same LSS that you used to start the session. This step is needed to later change the direction of the B volumes and restore the A volumes. Enter the pausegmir command at the dscli command prompt with the following parameters and variables:

dscli> pausegmir -dev IBM.2107-75ALA2P -quiet -lss 07 -session 1

The following represents an example of the output:

CMUC00165I pausegmir: Global Mirror for session 1 successfully paused.

See "Pausing Global Mirror processing" on page 589 for more information.

5. Suspend the B and C volume pairs. Because the site B volumes cannot be source volumes for Metro Mirror and Global Copy relationships, you must suspend the B volumes to C volumes so that B volumes to A volumes can be established.

This step stops all incoming write I/Os to the affected B and C volume pairs and helps prepare for a later resynchronization of the A volumes with the current operating B volumes.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 1200-125f:0700-075f

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

See "Pausing a Metro Mirror relationship" on page 572 for more information.

6. Establish paths between the local and intermediate sites that contain the B to A Metro Mirror volumes. Enter the mkpprcpath command at the dscli command prompt with the following parameters and variables:

dscli> mkpprcpath -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -remotewwnn 5005076303FFC550 -srclss 61 -tgtlss 63 -consistgrp I0102:I0031 I0002:I0102

The following represents an example of the output: CMUC00149I mkpprcpath: Remote Mirror and Copy path 61:63 successfully established.

See "Creating remote mirror and copy paths" on page 567 for more information.

7. Issue a failback command to the B volumes (with A volumes as secondaries): Host I/O processing continues uninterrupted to the B volumes as the A volumes are made current. This command copies the changes back to the A volumes that were made to the B volumes while hosts are running on the B volumes. (In a DS CLI environment, where the local and intermediate sites use different management consoles, you have to use a different DS CLI session for the management console of the B volumes at the intermediate site.) Enter the failbackpprc command at the dscli command prompt with the following parameters and variables:

```
dscli> failbackpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X
-type gcp 1200-125f:1a00-1a5f
```

The following represents an example of the output:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1200:1a00 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1201:1a01 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1202:1a02 successfully failed back.

See "Running a recovery failback operation" on page 577 for more information.

8. Wait for the copy process of the B to A volumes to reach full duplex (all out-of-sync tracks have completed copying). Host writes are no longer tracked. Monitor this activity by issuing queries to determine when the B to A volumes reach full duplex status (the number of out-of-sync tracks reaches zero). Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev IBM.2107-75ALA2P -remotedev -remotedev IBM.2107-75ALA2P -l -fmt default 1200-125f:1a00-1a5f

See "Viewing information about Metro Mirror relationships" on page 578 for more information.

- 9. Quiesce host I/O processing to the B volumes.
- **10. Issue a failover command to the A to B volumes.** This process ends the B to A volume relationships and establishes the A to B volume relationships. Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> failoverpprc -dev IBM.2107-130126X -remotedev IBM.2107-75ALA2P -type
mmir 1a00-1a5f:1200-125f
```

The following represents an example of the output:

```
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200
successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201
successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A02:1202
successfully reversed.
```

See "Running a recovery failover operation" on page 578 for more information.

11. Reestablish paths (that were disabled by the freeze operation) between the local site LSS and intermediate site LSS that contain the B to A Metro Mirror volume pairs. Enter the mkpprcpath command at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprcpath -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P
-remotewwnn 5005076303FFC550 -srclss 61 -tgtlss 63
-consistgrp I0102:I0031 I0002:I0102
```

The following represents an example of the output: CMUC00149I mkpprcpath: Remote Mirror and Copy path 61:63 successfully established.

See "Reestablishing remote mirror and copy paths (site A to site B)" on page 615 for more information.

12. Issue a failback command to the A to B volumes. This failback command completes the restore of the A to B volume relationship (the B volume becomes the target). The replication of the data starts immediately when the command is finished. Depending on how many tracks have changed during the disaster recovery test, resynchronization might take a long time.

Note: At this point, you can resume host I/O processing to the local site if optimizing host availability is critical. However, new host I/O that is copied to the A volumes at the local site is not fully protected by Global Mirror processing until the Global Mirror operation is restored in step 15.

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type mmir 1a00-1a5f:1200-125f

The following represents an example of the output:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A00:1200 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A01:1201 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A02:1202 successfully failed back.

See "Running a recovery failback operation" on page 577 for more information.

13. Reestablish Global Copy relationships between the B and C volumes with the Resync and Cascade options. When the failback operation has been done, Global Copy relationships can be re-created. Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760
-type gcp -mode nocp -cascade 1200-125f:0700-075f

The following represents an example of the output:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1200:0700 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1201:0701 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1202:0702 successfully created.

See "Creating a Global Copy relationship" on page 572 for more information.

14. Wait until the first pass of the Global Copy processing of the B and C volumes has completed. You can monitor this activity by querying the status of the B to C volume pairs in Global Copy relationships. Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-75ALA2P -l -fmt default 1200-125f:0700-075f

15. Resume Global Mirror. Now that the original infrastructure has been restored, you can resume the Global Mirror session. Enter the **resumegmir** command at the dscli command prompt with the following parameters and variables:

dscli> resumegmir -dev IBM.2107-75ALA2P -session 1 -lss 07

The following represents an example of the output:

CMUC00164I resumegmir: Global Mirror for session 1 successfully resumed.

See "Resuming Global Mirror processing" on page 589 for more information.

- **16. Resume host I/O processing to the A volumes.** Direct host I/O back to the A volumes in preparation for resuming host I/O on the A volumes.
- 17. Verify that consistency group are forming successfully. Enter the showgmir -metrics command at the dscli command prompt with the following parameters and variables:

dscli> showgmir -metrics 07

The following represents an example of the output:

ID	Total Failed CG Count	Total Succes- sful CG Count	Succes- sful CG Percen- tage	Failed CG after Last Success	Last Succes- sful CG Form Time	Coord. Time (milli- seconds)	CG Interval Time (sec- onds)
IBM.2107 -130165X /07	0	55	100	0	10/20/ 2005 11:38:25 MST	50	0

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
30	-	-	No Error	-	-	-	-

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
No Error	-	-	-	-	No Error	-	-

See "Querying Global Mirror processing" on page 588 for more information.

Failover and restore operations at the remote site during a planned outage

Use this process to run failover and restore operations (DS8000 only) to your remote (C) site during a planned outage. Because Global Mirror is not used at the intermediate site, E volumes were not included in this scenario.

Before you issue a failover operation to the remote site, ensure that data processing has completely stopped between the local and intermediate sites. If you fail to do so, data can be lost if you did not stop I/O to the A volumes at the local site before recovering at the remote site.

This scenario describes the steps in which a failover operation is done to move production from the local site to a remote site and then a failback operation is done when processing is ready to return to the local site. Assume that host I/O cannot be sent to the local site (Site A) in a Metro/Global Mirror configuration and it is not possible to run your systems using the B volumes at the intermediate site. You can switch operations to your remote site (Site C), which allows the processing of data to resume at the remote site. The Global Copy relationships between volumes at the intermediate and remote site are still operational. Global Mirror continues to operate between these two sites.

Note: For planned and unplanned outages at the local site and for certain disaster scenarios, GDPS HyperSwap can reset and restart systems using data from the B volumes following a failover operation.

GDPS HyperSwap does this transparently (without any system outage for systems running at the intermediate site) through the use of a *single* script statement.

Complete these steps for failover and restore operations at the remote site: (The steps in this scenario are examples.)

- 1. At the local site, ensure that data consistency is achieved between the A and B volume pairs. This process will help coordinate the A and B volumes consistency and allow consistent data to be copied to the C volumes at the remote site. You can use either one of the following methods:
 - Quiesce host I/O to the A volumes at the local site.
 - Freeze write activity to the Metro Mirror primary volumes.

If you quiesce I/O processing to the A volumes at the local site, continue to step 2 on page 637.

If you freeze write activity to the Metro Mirror primary volumes, complete the following steps:

a. Freeze updates to the A volumes in Metro Mirror relationships across the affected LSSs. This process ensures that the B volumes will be consistent at the time of the freeze. (One command per LSS is required.)

freezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07:12

The following example shows output that displays:

<code>CMUC00161W</code> freezepprc: Remote Mirror and Copy consistency group 07:12 successfully created.

As a result of the freeze action, the following processing occurs:

- I/O processing to the Metro Mirror volume pairs is temporarily queued during the time that updates are frozen.
- The volume pairs that are associated with the source and target LSSs are suspended. During this time, the storage unit collects data that is sent to the A volumes that are in Metro Mirror relationships.
- The established paths between the LSS pairs are disabled.
- b. If wanted, you can view the state of the pair status at the local site after the **freezepprc** command has been processed. Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> lspprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -fmt default 0700-075f
```

The following example shows output that displays:

Notes:

- 1) The command example uses the command parameter **-fmt default**. This command parameter specifies that the output be set to a space-separated plain text table.
- **2)** The following table format is presented for clarity. The actual report is not displayed in this format.
- **3**) The report example represents the information that is reported on when you do not specify the **-1** parameter.

ID	State	Reason	Туре	Source- LSS	Timeout (secs)	Critical Mode	First Pass Status
0700:1200	Suspend- ed	Freeze	Metro Mirror	07	unknown	Disabled	Invalid
0701:1201	Suspend- ed	Freeze	Metro Mirror	07	unknown	Disabled	Invalid
0702:1202	Suspend- ed	Freeze	Metro Mirror	07	unknown	Disabled	Invalid

c. Resume operations following a freeze.

Issue the **unfreezepprc** command to allow I/O activity to resume for the specified volume pairs. Enter the **unfreezepprc** command at the dscli command prompt with the following parameters and variables:

Note: This activity is sometimes referred to as a *thaw* operation. dscli> unfreezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07:12 The following example shows output that displays: CMUC00198I unfreezepprc: Remote Mirror and Copy pair 07:12 successfully thawed.

2. Verify that the last data from the local site has been included in a Global Mirror consistency group. Monitor this activity to determine when at least two consistency groups have formed since I/O processing was quiesced or freeze commands were issued to the local site. The "Total Successful CG Count" field from the query output displays this information.

At this point, data on the B, C, and D volumes is consistent. Enter the **showgmir** -**metrics** command at the dscli command prompt with the following parameters and variables:

dscli> showgmir -metrics 10

ID	Total Failed CG Count	Total Succes- sful CG Count	Succes- sful CG Percen- tage	Failed CG after Last Success	Last Succes- sful CG Form Time	Coord. Time (milli- seconds)	CG Interval Time (sec- onds)
IBM.2107 -75ALA 2P/10	23	139	85	0	02/20/ 2006 11:33:56 MST	50	0

The following example shows output that displays:

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
30	IBM.2107 -75ALA 2P	0x12	Error	Session or Session Members not in Correct State	Global Mirror Run in Progress	IBM.2107 -75ALA 2P	Not Available

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	Max Drain Time Exceeded	Drain in Progress	IBM.2107 -75ALA 2P	Not Available	Error	Max Drain Time Exceeded	Drain in Progress

3. End the Global Mirror session between the B and C volume pairs.

Enter the **rmgmir** command at the dscli command prompt with the following parameters and variables:

dscli> rmgmir -quiet -lss 10 -session 31

The following example shows output that displays:

CMUC00165I rmgmir: Global Mirror for session 31 successfully stopped.

See "Ending Global Mirror processing (script mode)" on page 590 or "Ending Global Mirror processing (no script)" on page 591 for more information.

4. Verify that the Global Mirror session has ended. Consistency groups will not be forming when Global Mirror processing is stopped.

Enter the **showgmir** command at the dscli command prompt with the following parameters and variables:

dscli> showgmir -metrics 10

The following example shows output that displays:

ID	Total Failed CG Count	Total Succes- sful CG Count	Succes- sful CG Percen- tage	Failed CG after Last Success	Last Succes- sful CG Form Time	Coord. Time (milli- seconds)	CG Interval Time (sec- onds)
IBM.2107 -75ALA 2P/10	414	148054	99	0	03/20/ 2006 15:33:56 MST	50	0

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
30	IBM.2107 -75ALA 2P	0x12	Error	Session or Session Members not in Correct State	Global Mirror Run in Progress	IBM.2107 -75ALA 2P	Not Available

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	Max Drain Time Exceeded	Drain in Progress	IBM.2107 -75ALA 2P	Not Available	Error	Max Drain Time Exceeded	Drain in Progress

5. Delete the relationships between the B and C volume pairs between the intermediate and remote sites. This prepares for reversing the direction of the volume pair from the remote site to the intermediate site. The cascaded relationship ends as well.

Note: When the relationships between the B and C volumes are deleted, the cascade parameter is disabled for the B volumes and the B volumes are no longer detected as being in cascaded relationships.

Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -quiet -dev IBM.IBM.2107-75ALA2P -remotedev IBM.2107-1831760
1200-125f:0700-075f
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1200:0700 relationship
successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1201:0701 relationship
successfully withdrawn.
CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1202:0702 relationship
successfully withdrawn.

6. Issue a failover command to the B and A volume pairs, with the Cascade option. With this process, updates are collected using the change recording feature, which allows for the resynchronization of the B and A volumes. Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failoverpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -type
gcp -cascade 1200-125f:1a00-1a5f

The resulting output (a shortened version) displays:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1200:1A00 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1201:1A01 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1202:1A02 successfully reversed.

7. **Create Global Copy relationships using the C and B volume pairs.** Specify the NOCOPY option. You can specify the NOCOPY option with the following command because the B and C volumes contain exact copies of data. Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P -type gcp -mode nocp 0700-075f:1200-125f

The following example shows output that displays:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0700:1200 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0701:1201 successfully created.

- 8. Start I/O processing at the remote site. Continue in this mode until production is ready to return to the local site.
- 9. When you are ready to return production to the local site, quiesce I/O processing at the remote site. This process is used to begin the transition back host I/O to the A volumes.
- 10. Wait for the number of out-of-sync tracks on the C and B volume to reach zero. You can monitor this activity by querying the status of the C and B volumes. As soon as the number of out-of-sync tracks reaches zero, all data has been copied and the data on the C and B volumes is equal. All updates that are needed to resynchronize the A volumes are recorded at the B volumes. Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P -l -fmt default 0700-075f

ID	State	Reason	Туре	Out Of Sync Tracks	Tgt Read	Src Cascade
0700:1200	Copy Pending	-	Global Copy	0	Disabled	Disabled
0701:1201	Copy Pending	-	Global Copy	0	Disabled	Disabled
0702:1202	Copy Pending	-	Global Copy	0	Disabled	Disabled

The following example shows output that displays:

Tgt Cascade	Date Sus pended	Source LSS	Timeout (secs)	Crit Mode	First Pass Status	Incre mental Resync	Tgt Write
Invalid	-	07	Unknown	Disabled	True	Enabled	Enabled
Invalid	-	07	Unknown	Disabled	True	Enabled	Enabled
Invalid	-	07	Unknown	Disabled	True	Enabled	Enabled

11. Reestablish paths (that were disabled by the freeze operation) between the local site LSS and intermediate site LSS that contain the B to A Metro Mirror volume pairs. Enter the mkpprcpath command at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprcpath -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P
-remotewwnn 5005076303FFC550 -srclss 61 -tgtlss 63 -consistgrp
I0102:I0031 I0002:I0102
```

The following example shows output that displays:

CMUC00149I mkpprcpath: Remote Mirror and Copy path 61:63 successfully established.

See "Creating remote mirror and copy paths" on page 567 for more information.

12. Issue a failback command to the B to A volume pairs. This command copies the changes back to the A volumes that were made to the B volumes while hosts were running on the B volumes. The A volumes are now synchronized with the B volumes. (In a DS CLI environment, where the local and intermediate sites use different management consoles, you have to use a different DS CLI session for the management console of the B volumes at the intermediate site.) Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -type gcp 1200-125f:1a00-1a5f

The following example shows output that displays:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1200:1A00 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1201:1A01 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1202:1A02 successfully failed back.

See "Running a recovery failback operation" on page 577 for more information.

- 13. Wait for the copy process of the B and A volume pairs to reach full duplex (all out-of-sync tracks have completed copying). You can monitor this activity by querying the status of the B and A volumes. As soon as the number of out-of-sync tracks reaches zero, all data has been copied and the data on the B and A volumes is equal. At this point, the data on volumes A, B, and C is equal. Enter the lspprc command at the dscli command prompt with the following parameters and variables: dscli>lspprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -1 -fmt default 1200-125f:1a00-1a5f
- 14. Delete the Global Copy relationships between the C and B volume pairs between the intermediate and remote sites. Deleting the Global Copy relationships between the C to B volume pairs prepares for restoring to the original Global Copy relationships between the B to C volume pairs. Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -quiet -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P 0700-075f:1200-125f

The following example shows output that displays:

CMUC00155I rmpprc: Remote Mirror and Copy volume pair 0700:1200 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair 0701:1201 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair 0702:1202 relationship successfully withdrawn.

15. Issue a failover command to the A and B volume pairs. This process ends the Metro Mirror relationships between the B and A volumes and establishes the Metro Mirror relationships between the A and B volumes. Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failoverpprc -dev IBM.2107-130165X -remotedev IBM.IBM.2107-75ALA2P -type mmir la00-la5f:1200-125f

The following example shows output that displays:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A02:1202 successfully reversed.

16. Reestablish paths (that were disabled by the freeze operation) between the local site LSS and the intermediate site LSS that contain the B to A Metro Mirror volume pairs. Enter the mkpprcpath command at the dscli command prompt with the following parameters and variables:

dscli> mkpprcpath -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -remotewwnn 5005076303FFC550 -srclss 61 -tgtlss 63 -consistgrp I0102:I0031 I0002:I0102

The following example shows output that displays:

CMUC00149I mkpprcpath: Remote Mirror and Copy path 61:63 successfully established.

See "Creating remote mirror and copy paths" on page 567 for more information.

17. Issue a failback command to the A to B volumes. This command copies the changes back to the A volumes that were made to the B volumes in Metro Mirror relationships while hosts were running on the B volumes. The A volumes are now synchronized with the B volumes. (In a DS CLI environment, where the local and intermediate sites use different management consoles, you have to use a different DS CLI session for the management console of the B volumes at the intermediate site.) Enter the failbackpprc command at the dscli command prompt with the following parameters and variables:

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type mmir 1a00-1a5f:1200-125f

The following example shows output that displays:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A00:1200 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A01:1201 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A02:1202 successfully failed back.

See "Running a recovery failback operation" on page 577 for more information.

18. Reestablish the B to C volume pairs in Global Copy relationships. Specify the NOCOPY option and the Cascade options. Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 -type gcp -mode nocp -cascade 1200-125f:0700-075f

The following example shows output that displays:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1200:0700 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1201:0701 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1202:0702 successfully created.

See "Creating a Global Copy relationship" on page 572 for more information.

19. Use FlashCopy to create a copy of C source volumes to the D target volumes. Enter the **mkflash** command at the dscli command prompt with the following parameters and variables. This creates a backup copy of the consistency group.

dscli> mkflash -dev IBM.2107-1831760 -tgtinhibit -record -persist -nocp 1300-125f:1900-195f

The following example shows output that displays: CMUC00137I mkflash: FlashCopy pair 1300:1900 successfully created. CMUC00137I mkflash: FlashCopy pair 1301:1901 successfully created.

CMUC001371 mkflash: FlashCopy pair 1301:1901 Successfully created. CMUC00137I mkflash: FlashCopy pair 1302:1902 successfully created.

See "Creating FlashCopy relationships (Global Mirror setup)" on page 594 for more information.

20. Resume Global Mirror processing. Enter the **resumegmir** command at the dscli command prompt with the following parameters and variables:

dscli> resumegmir -dev IBM.2107-75ALA2P -session 31 -lss

The following example shows output that displays:

CMUC00164I resumegmir: Global Mirror for session 31 successfully resumed.

See "Resuming Global Mirror processing" on page 589 for more information.

21. Resume host I/O processing to the A volumes.

Failover and restore operations at the remote site during an unplanned outage

Use this process to run failover and restore operations (DS8000 only) at your remote (C) site during an unplanned outage, using E volumes at the intermediate site.

If possible, before you issue a failover operation to the remote site, ensure that data processing has completely stopped at the local and intermediate sites. If you fail to do so, data can be lost if you did not quiesce I/O processing to the local site before recovering at the remote site.

For this scenario, assume that host I/O processing is being sent to the local site in a Metro/Global Mirror configuration. A failure occurs at the local (A) site and it is not possible to run your systems using the B volumes at the intermediate site. You can switch operations to your remote (C) site, which allows the processing of data to resume at site C. This process is known as a failover recovery. The Global Copy relationship between volumes at the intermediate and remote site is still operational. Global Mirror continues to operate between these two sites.

Note: The steps to run a failover operation to the remote site (C) can be done using a *single* GDPS automation script running in a GDPS Global Mirror remote controlling system. GDPS provides a recovery check function that determines the state of the volumes *before* the necessary actions to run a recovery function are done. This process alerts the user to fix the required problem before running the actual recovery, which can reduce the time required to complete it.

Complete the following steps after a failure has been detected at the local site. (The steps in this scenario are examples.)

- 1. If the local site was not completely destroyed, it is essential that data from any surviving A and B volume pairs be copied and a consistent copy be achieved at the remote site. If possible and you are able to freeze write activity to the Metro Mirror primary volumes, complete the following steps:
 - a. Freeze updates to the A volumes in Metro Mirror relationships across the affected LSSs. This process ensures that the B volumes are consistent at the time of the freeze. (One command per LSS is required.) Enter the **freezepprc** command at the dscli command prompt with the following parameters and variables:

dscli> freezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07:12

The following represents an example of the output:

<code>CMUC00161W</code> freezepprc: Remote Mirror and Copy consistency group 07:12 successfully created.

As a result of the freeze action, the following processing occurs:

- The established paths between the LSS pairs are deleted.
- The volume pairs that are associated with the source and target LSSs are suspended. During this time, the storage unit collects data that is sent to the A volumes that are in Metro Mirror relationships.
- I/O processing to the Metro Mirror volume pairs is temporarily queued during the time that updates are frozen.
- b. Resume operations following a freeze.

Issue the **unfreezepprc** command to allow I/O activity to resume for the specified volume pairs. Enter the **unfreezepprc** command at the dscli command prompt with the following parameters and variables:

Note: This activity is sometimes referred to as a *thaw* operation. dscli> unfreezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07:12 The following represents an example of the output:

CMUC00198I unfreezepprc: Remote Mirror and Copy pair 07:12 successfully thawed.

2. Verify that the last data from the local site has been included in a Global Mirror consistency group. Monitor this activity by querying the B and C volumes to determine when at least two successful consistency groups have formed. The "Total Successful CG Count" field from the query output displays this information.

Note: When you use the **showgmir** command with the **-metrics** parameter, you can monitor the progress of the consistency group formation. When Global Mirror is running, the number of consistency groups is steadily growing each time you issue the **showgmir** command.

Enter the **showgmir** -**metrics** command at the dscli command prompt with the following parameters and variables:

dscli> showgmir -metrics 10

The following represents an example of the output:

ID	Total Failed CG Count	Total Succes- sful CG Count	Succes- sful CG Percen- tage	Failed CG after Last Suc- cess	Last Succes- sful CG Form Time	Coord. Time (milli- seconds)	CG Inter- val Time (sec- onds)
IBM. 2107- 75ALA 2P/10	23	139	85	0	02/20/ 2006 11:33:56 MST	50	0

Max CG Drain Time (sec- onds)	First Failure Control Unit	First Failure LSS	First Failure Status	First Failure Reason	First Failure Master State	Last Failure Control Unit	Last Failure LSS
30	IBM. 2107- 75ALA 2P	0x12	Error	Session or Session Members not in Correct State	Global Mirror Run in Progress	IBM.2107- 75ALA 2P	Not Available

Last Failure Status	Last Failure Reason	Last Failure Master State	Prev- ious Failure Control Unit	Prev- ious Failure LSS	Prev- ious Failure Status	Prev- ious Failure Reason	Prev- ious Failure Master State
Error	Max Drain Time Exceed- ed	Drain in Prog- ress	IBM.2107 -75ALA 2P	Not Avail- able	Error	Max Drain Time Exceed- ed	Drain in Prog- ress

3. Stop the Global Mirror session from which the B and C volume pairs are included. Enter the rmgmir command at the dscli command prompt with the following parameters and variables:

dscli> rmgmir -quiet -lss 10 -session 31

The following represents an example of the output:

CMUC00165I rmgmir: Global Mirror for session 31

successfully stopped.

See "Ending Global Mirror processing (script mode)" on page 590 or "Ending Global Mirror processing (no script)" on page 591 for more information.

4. Verify that the Global Mirror session has ended. Consistency groups do not form when Global Mirror processing is stopped.

See "Querying Global Mirror processing" on page 588 for more information.

Enter the **showgmir** command at the dscli command prompt with the following parameters and variables:

dscli> showgmir 10

The following represents an example of the output:

ID	Master Count	Master Session ID	Copy State	Fatal Reason	CG Inter- val (sec- onds)	Coord. Time (milli- sec- onds)	CG Drain Time (sec- onds)	Current Time
IBM. 2107- 75ALA 2P/10	-	-	-	-	-	-	-	-

CG Time	Succes- sful CG Percen- tage	Flash- Copy Sequ- ence Number	Master ID	Subor- dinate Count	Master/ Subordi- nate Assoc.
-	-	-	-	-	-

5. Delete the Global Copy relationships between the B and C volume pairs at the intermediate and remote sites. When the relationships between the B and C volumes are deleted, the cascade parameter is disabled for the B volumes and the B volumes are no longer detected as being in cascaded relationships. Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> rmpprc -quiet -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 1200-125f:0700-075f
```

The following represents an example of the output:

CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1202:0702 relationship successfully withdrawn.

See "Deleting a Metro Mirror relationship" on page 573 for more information.

6. **Issue a failover command to the B volumes with the Cascade option.** With this process, updates are collected using the change recording feature, which allows the later resynchronization of the B to A volumes. Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failoverpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -type gcp -cascade 1200-125f:1a00-1a5f

The following represents an example of the output:

```
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1200:1A00
successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1201:1A01
successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1202:1A02
successfully reversed.
```

See "Running a recovery failover operation" on page 578 for more information.

7. **Create Global Copy relationships using the C and B volume pairs.** Specify the NOCOPY option. Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

Note: You can specify the NOCOPY option with the following commands because the B and C volume pairs contain exact copies of data.

dscli> mkpprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P -type gcp -mode nocp 0700-075f:1200-125f

The following represents an example of the output:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0700:1200 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0701:1201 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 0702:1202 successfully created.

See "Creating a Global Copy relationship" on page 572 for more information.

8. Use FlashCopy to create a copy of B source volumes to E target volumes. Specify the following options: Persistent and Start Change Recording. This creates a backup copy of the consistency group. Enter the **mkflash** command at the dscli command prompt with the following parameters and variables:

```
dscli> mkflash -dev IBM.2107-75ALA2P -tgtinhibit -record -persist -nocp 1200-125f:1900-195f
```

The following represents an example of the output:

CMUC00137I mkflash: FlashCopy pair 1200:1900 successfully created. CMUC00137I mkflash: FlashCopy pair 1201:1901 successfully created. CMUC00137I mkflash: FlashCopy pair 1202:1902 successfully created.

See "Creating FlashCopy relationships (Global Mirror setup)" on page 594 for more information.

9. Create a Global Mirror session using the C volumes. Enter the mksession command at the dscli command prompt with the following parameters and variables:

```
dscli> mksession -lss 07 1
```

The following represents an example of the output: CMUC00145I mksession: Session 1 opened successfully.

See "Creating the Global Mirror session" on page 595 for more information.

10. Start the Global Mirror session from which the C, B and E volumes are included. Enter the mkgmir command at the dscli command prompt with the following parameters and variables: dscli> mkgmir -lss 07 -session 1

The following represents an example of the output:

CMUC00162I mkgmir: Global Mirror for session 1 successfully started.

See "Starting Global Mirror processing" on page 590 for more information.

- 11. Verify that the Global Mirror session has started. Enter the showgmir command at the dscli command prompt with the following parameters and variables:
 - dscli> showgmir 07

CG Coord. CG Inter-Time Drain Master val (milli-Time Master Session Fatal (sec-(sec-Current Copy sec ID ID Reason Time Count State onds) onds) onds) IBM. 1 0×01 Not Fatal 0 50 30 02/20/ Running 2107-2006 75ALA 11:37:40 2P/07 MST

The following represents an example of the output:

CG Time	Succes- sful CG Percen- tage	Flash- Copy Sequ- ence Number	Master ID	Subor- dinate Count	Master/ Subordi- nate Assoc.
02/20/ 2006 11:37:40 MST	100	0x4357E- 3F4	IBM. 2107- 75ALA 2P/07	0	-

12. Allow the I/O to run and monitor the formation of the consistency groups. Enter the showgmir command at the dscli command prompt with the following parameters and variables:

dscli> showgmir 07

The following represents an example of the output:

ID	Master Count	Master Session ID	Copy State	Fatal Reason	CG Inter- val (sec- onds)	Coord. Time (milli- sec- onds)	CG Drain Time (sec- onds)	Current Time
IBM. 2107- 75ALA 2P/07	1	0x01	Running	Not Fatal	0	50	30	02/20/ 2006 11:37:40 MST

CG Time	Succes- sful CG Percen- tage	Flash- Copy Sequ- ence Number	Master ID	Subor- dinate Count	Master/ Subordi- nate Assoc.
02/20/ 2006 11:37:40 MST	100	0x4357E3- F4	IBM. 2107- 75ALA 2P/07	0	-

13. When the local site is ready to return, issue a failback command to the B and A volumes. Before the applications are started at the local site, data at the local site has to be copied from the intermediate site. Issue the failbackpprc command to start copying data from the B volumes at the intermediate site to the A volumes at the local site while hosts are running on the B volumes. When all data is copied, the A volumes are synchronized with the B volumes.

Note: In a DS CLI environment, where the local and intermediate sites use different management consoles, you have to use a different DS CLI session for the management console of the B volumes at the intermediate site.

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type gcp -cascade 1200-125f:1a00-1a5f

The following represents an example of the output:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1200:1A00 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1201:1A01 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1202:1A02 successfully failed back.

See "Running a recovery failback operation" on page 577 for more information.

14. Wait for the copy operation of the B and A volumes to reach full duplex status (all out-of-sync tracks have completed copying). You can monitor this activity by querying the status of the B and A volume pairs. Enter the lspprc command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -l -fmt default 1200-125f

The following represents an example of the output:

ID	State	Reason	Туре	Out Of Sync Tracks	Tgt Read	Src Cascade
1200:1a00	Copy Pending	-	Global Copy	0	Disabled	Enabled
1201:1a01	Copy Pending	-	Global Copy	0	Disabled	Enabled
1201:1a02	Copy Pending	-	Global Copy	0	Disabled	Enabled

Tgt Cascade	Date Sus pended	Source LSS	Timeout (secs)	Crit Mode	First Pass Status	Incre mental Resync	Tgt Write
Invalid	-	12	Unknown	Disabled	True	Enabled	Enabled
Invalid	-	12	Unknown	Disabled	True	Enabled	Enabled
Invalid	-	12	Unknown	Disabled	True	Enabled	Enabled

15. End I/O processing to the C volumes. Enter the **rmgmir** command at the dscli command prompt with the following parameters and variables:

dscli> rmgmir -quiet -lss 07 -session 1

The following represents an example of the output:

CMUC00165I rmgmir: Global Mirror for session 1 successfully stopped.

See "Ending Global Mirror processing (script mode)" on page 590 or "Ending Global Mirror processing (no script)" on page 591 for more information.

16. Verify that at least two consistency groups have formed. Assuming that the consistency groups formed successfully, the A, B, C, and E volumes contain consistent data. (Data at the remote site is consistent to the last successful consistency group formed by the master storage unit.)See "Querying Global Mirror processing" on page 588 for more information.

Enter the **showgmir** -**metrics** command at the dscli command prompt with the following parameters and variables:

dscli> showgmir -metrics 07

The following represents an example of the output:

ID	Total Failed CG Count	Total Succes- sful CG Count	Succes- sful CG Percent- age	Failed CG after Last Success	Last Succes- sful CG Form Time	Coord. Time (milli- sec- onds)	CG Interval Time (sec- onds)
IBM.2107 -75ALA 2P/10	0	55	100	0	02/20/ 2006 11:38:25 MST	50	0

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
30	-	-	No Error	-	-	-	-

Last Failure Status	Last Failure Reason	Last Failure Master State	Prev- ious Failure Control Unit	Prev- ious Failure LSS	Previous Failure Status	Previous Failure Reason	Prev- ious Failure Master State
No Error	-	-	-	-	No Error	-	-

17. End the Global Mirror session between the C, B, and E volumes. Enter the rmgmir command at the dscli command prompt with the following parameters and variables:

dscli> rmgmir -quiet -lss 07 -session 1

The resulting output is displayed:

 $\verb|CMUC00165I rmgmir: Global Mirror for session 1 successfully stopped||$

See "Ending Global Mirror processing (script mode)" on page 590 or "Ending Global Mirror processing (no script)" on page 591 for more information.

18. Verify that the Global Mirror session that includes the C, B, and E volumes has stopped. Enter the showgmir command at the dscli command prompt with the following parameters and variables. showgmir 07

ID	Total Failed CG Count	Total Succes- sful CG Count	Succes- sful CG Percen- tage	Failed CG after Last Success	Last Succes- sful CG Form Time	Coord. Time (milli- sec- onds)	CG Interval Time (sec- onds)
IBM.2107 -75ALA 2P/10	23	139	85	0	02/20/ 2006 11:33:56 MST	50	0

The following represents an example of the output:

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
30	IBM.2107 -75ALA 2P	0x12	Error	Session or Session Members not in Correct State	Global Mirror Run in Progress	IBM.2107 -75ALA 2P	Not Available

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	Max Drain Time Exceeded	Drain in Progress	IBM.2107 -75ALA 2P	Not Available	Error	Max Drain Time Exceeded	Drain in Progress

See "Querying Global Mirror processing" on page 588 for more information.

19. At the remote site, remove the C volumes (or Global Copy secondary volumes) from the Global Mirror session that includes the C, B, and E volumes. Enter the chsession command at the dscli command prompt with the following parameters and variables:

```
dscli> chsession -dev BM.2107-75ALA2P -action remove -volume 1200-125f
-lss 07 1
The resulting output is displayed:
```

CMUC00147I chsession: Session 1 successfully modified.

See "Removing volumes from a session (Global Mirror)" on page 596 for more information.

20. Delete the Global Copy relationships between the C to B volumes between the intermediate and remote sites. Deleting the Global Copy relationships between the C to B volume pairs prepares for restoring to the original Global Copy relationships between the B to C volume pairs. The cascaded relationship ends, as well.

Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -quiet -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P 0700-075f:1200-125f CMUC00155I rmpprc: Remote Mirror and Copy volume pair 0700:1200 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair 0701:1201 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair 0702:1202 relationship successfully withdrawn.

See "Deleting a Metro Mirror relationship" on page 573 for more information.

21. Issue a failover command to the A to B volumes. This process ends the Metro Mirror relationships between the B and A volumes and establishes the Metro Mirror relationships between the A and B volume pairs. Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

failoverpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type mmir 1a00-1a5f:1200-125f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A02:1202 successfully reversed.

22. Reestablish paths that were disabled by the freeze operation between the local site LSS and intermediate site LSS that contain the B to A Metro Mirror volume pairs. Enter the mkpprcpath command at the dscli command prompt with the following parameters and variables:

dscli> mkpprcpath -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -remotewwnn 5005076303FFC550 -srclss 61 -tgtlss 63 -consistgrp I0102:I0031 I0002:I0102 The resulting output is displayed:

CMUC00149I mkpprcpath: Remote Mirror and Copy path 61:63 successfully established.

See "Reestablishing remote mirror and copy paths (site A to site B)" on page 615 for more information.

23. Issue a failback command to the A and B volumes. This command copies the changes back to the A volumes that were made to the B volumes in Metro Mirror relationships while hosts were running on the B volumes. The A volumes are now synchronized with the B volumes. (In a DS CLI environment, where the local and intermediate sites use different management consoles, you have to

use a different DS CLI session for the management console of the B volumes at the intermediate site.) Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type mmir 1a00-1a5f:1200-125f

The following represents an example of the output:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A00:1200 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A01:1201 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A02:1202 successfully failed back.

24. Establish the B and C volume pairs in Global Copy relationships. Enter the mkpprc command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 -type gcp -mode nocp -cascade 1200-125f:0700-075f

The following represents an example of the output:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1200:0700 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1201:0701 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1202:0702 successfully created

See "Creating a Global Copy relationship" on page 572 for more information.

25. Optionally, you can issue a FlashCopy operation to create a backup copy of all the C, B, and E volumes from which the last consistency group was created. If you need to preserve data from the set of volumes (or consistency group) that was created using the E volumes, allow the background copy from the FlashCopy process to complete before you continue to the next step, which describes removing the FlashCopy relationship between the B to E volume pairs.

Enter the **mkflash** command at the dscli command prompt with the following parameters and variables:

dscli> mkflash -dev IBM.2107-75ALA2P -tgtinhibit -record -persist -nocp 1200-125f:1900-195f

The following represents an example of the output:

CMUC00137I mkflash: FlashCopy pair 1200:1900 successfully created. CMUC00137I mkflash: FlashCopy pair 1201:1901 successfully created.

See "Creating FlashCopy relationships (Global Mirror setup)" on page 594 for more information.

26. Delete the FlashCopy relationship between the B and E volume pairs to end the relationship at the intermediate site. Enter the rmflash command at the dscli command prompt with the following parameters and variables:

dscli> rmflash -dev IBM.2107-75ALA2P -quiet 1200-125f:1900-195f

The following represents an example of the output:

CMUC00140I rmflash: FlashCopy pair 1200:1900 successfully removed. CMUC00140I rmflash: FlashCopy pair 1201:1901 successfully removed.

See "Removing FlashCopy relationships" on page 598 for more information.

27. Resume Global Mirror at the intermediate site. This starts Global Mirror processing for the B, C , and D volumes.

Enter the **resumegmir** command at the dscli command prompt with the following parameters and variables:

dscli> resumegmir -dev IBM.2107-75ALA2P -session 10 -lss 31

The resulting output is displayed:

CMUC00164I resumegmir: Global Mirror for session 10 successfully resumed.

See "Resuming Global Mirror processing" on page 589 for more information.

28. Resume I/O on A volumes.

29. Verify that consistency groups are forming successfully.

Enter the **showgmir** -metrics command at the dscli command prompt with the following parameters and variables:

dscli> showgmir -metrics 10

The following represents an example of the output:

ID	Total Failed CG Count	Total Succes- sful CG Count	Succes- sful CG Percen- tage	Failed CG after Last Success	Last Succes- sful CG Form Time	Coord. Time (milli- sec- onds)	CG Inter- val Time (sec- onds)
IBM.2107 -75ALA 2P/10	1	39	97	0	02/20/ 2006 11:33:56 MST	50	0

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
30	IBM.2107 -75ALA 2P	0x12	Error	Session or Session Members not in Correct State	Global Mirror Run in Progress	IBM.2107 -75ALA 2P	0x12

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	Session or Session Members not in Correct State	Global Mirror Run in Progress	-	-	No Error	-	Drain in Progress

See "Querying Global Mirror processing" on page 588 for more information.

Using forced failover and failback during a planned Metro/Global Mirror outage

Use this process to run failover operations (DS8000 only) from the local (A) site to the intermediate (B) site during a planned outage.

The following assumptions apply to your 3-site Metro/Global Mirror configuration:

- You used incremental resynchronization to establish the relationship between your site A and site B volumes and between your site A and site C volumes.
- You established a Global Mirror session at the local site. This means that Fibre Channel paths were established between all Global Mirror source and target pairs and between the master and subordinate storage units.

The command examples use the following site identifiers:

- Site A is identified as 2107-130165X
- Site B is identified as 2107-75ALA2P
- Site C is identified as 2107-183176O

This process uses forced failover and forced failback processing to establish a relationship between the C and A volumes without verification that this relationship already existed. A **-force** parameter has been added to the **failoverpprc** and **failbackpprc** commands to accommodate this processing.

Attention: Use the **-force** parameter only as directed. This parameter can cause severe damage to your data if it is misused. Contact IBM Support before you attempt to use the **-force** parameter if your outage situation is outside the boundaries of this example.

Complete the following steps to failover and restore operations to the intermediate (B) site.

1. At the local site, ensure that data consistency is achieved between the site A and site B volumes. You can use freeze and unfreeze commands that are supported using external automation software to create data consistency to multiple Metro Mirror volume pairs.

To freeze write activity to Metro Mirror primary volumes, complete the following steps:

a. Freeze updates to the A volumes in Metro Mirror relationships across the affected LSSs. This ensures that the B volumes will be consistent at the time of the freeze process. (One command per LSS is required.)

Enter the **freezepprc** command at the dscli command prompt with the following parameters and variables:

dscli> freezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07-12

The following represents an example of the output:

<code>CMUC00161W</code> freezepprc: Remote Mirror and Copy consistency group 07:12 successfully created.

As a result of the freeze action, the following processing occurs:

- The established Remote Mirror and Copy paths between the LSS pairs are deleted.
- The volume pairs that are associated with the source and target LSSs are suspended. During this time, the storage unit collects data that is sent to the A volumes in Metro Mirror relationships.
- I/O to the Metro Mirror volume pairs is temporarily queued.
- b. **Resume operations following a freeze.** This operation—also called a thaw operation—allows I/O processing to resume for the specified volume pairs.

Enter the **unfreezepprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> unfreezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P
07:12
```

```
The following represents an example of the output:
CMUC00198I unfreezepprc: Remote Mirror and Copy pair 07:12
successfully thawed.
```

2. Create a relationship from the C volumes to the A volumes, using the -force and -cascade parameters. No validation is done at site C to determine that site C is a secondary of site A.

Note: For this step to succeed you must ensure that the Remote Mirror and Copy paths between all Global Mirror source and target pairs and between the Master and subordinate storage units have been created.

Enter the failoverpprc command at the dscli prompt with the following parameters and variables:

dscli> failoverpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp -cascade -force SourceVolumeID:TargetVolumeID

Example

dscli> failoverpprc -dev IBM.2107-1831760 -remotedev IBM.2107-130165X -type gcp -cascade -force 1200-125f:1A00-1A5f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1200:1A00 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1201:1A01 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1202:1A02 successfully reversed.

3. End the Metro Mirror relationship between the A to B volumes at the B volumes intermediate site.

Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> rmpprc -dev storage_image_ID -remotedev storage_image_ID
-at tgt -unconditional -quiet TargetVolumeID
```

Example

dscli> rmpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -at tgt -unconditional -quiet 1200-125f

4. **Redirect host I/O processing to the B volumes.** Changes are recorded on the B volumes until the A volumes can be resynchronized with the B volumes. Also, Global Mirror continues to operate from site B to site C.

Note: You can run in this configuration until the A site has recovered and you want to restore operations there. Begin the next step after the A volumes have been recovered and you're still in production on the B volumes.

5. Copy changes from site C back to site A, using the -force parameter. Host I/O processing continues uninterrupted to the B volumes while the A volumes are made current. (The data is still flowing from B to C, so any changes made to B are being transferred to C and therefore will get from C to A.) This command copies the changes back to the A volumes that were made to the B volumes while hosts were running on the A volumes. (In a DS CLI environment, where the local and remote sites are not using the same management console, you have to use the management console of the remote site.)

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp -force SourceVolumeID:TargetVolumeID

Example

```
dscli> failbackpprc -dev IBM.2107-1831760 -remotedev IBM.2107-130165X
-type gcp -force 1200-125f:1A00-1A5f
```

The following represents an example of the output:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1200:1A00 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1201:1A01 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1202:1A02 successfully failed back.

6. Wait for the first pass copy to complete from site C to site A. Issue the lspprc command if you want to monitor this activity and determine when the first pass status changes to "True."

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID
-l SourceVolumeID:TargetVolumeID

Example

dscli> lspprc -dev IBM.2107-1831760 -remotedev IBM.2107-75130165X -l 1200-125f:1A00-1A5f

The following represents the first two lines of the report generated by the **lspprc** command:

ID	State	Reason	Туре	Source- LSS	Time- out (secs)	Critical Mode	First Pass Status
IBM.2107- 183176O /2101: IBM.2107- 130165X /2101	Copy Pend- ing	-	Global Copy	IBM.2107- 130165X /20	300	Disabled	True
IBM.2107- 183176O /2100: IBM.2107- 130165X /2100	Copy Pend- ing	-	Global Copy	IBM.2107- 130165X /20	300	Disabled	True

7. Modify Global Copy relationships between the B and C volume pairs. Specify the NOCOPY option and initiate incremental resynchronization without initialization.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID -type gcp -incrementalresync enablenoinit -mode nocp SourceVolumeID:TargetVolumeID

Example

dscli> mkpprc -dev IBM.2107-75ALA2P -remotedev IBM2107-1831760 -type gcp -incrementalresync enablenoinit -mode nocp 1200-125f:1A00-1A5f

The following represents the first two lines of the output:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1200:1A00 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1201:1A01 successfully created.

8. Begin the process to return production to site A. First, the Global Mirror session at site B must be stopped.

Enter the **rmgmir** command at the dscli command prompt with the following parameters and variables:

dscli> rmgmir -dev storage_image_ID -lss LSS_ID -session session_ID Master Control Path LSS ID:Subordinate Control Path LSS ID

Example

dscli> rmgmir -dev IBM.2107-75ALA2P -quiet -lss 07 -session 1

The following represents an example of the output:

CMUC00165I rmgmir: Global Mirror for session 1 successfully terminated.

9. Verify that the Global Mirror session has ended.

Enter the **showgmir** command at the dscli command prompt with the following parameters and variables:

dscli> showgmir -dev storage_image_ID LSS_ID

Example

dscli> showgmir -dev IBM.2107-75ALA2P 10

In the resulting report, the output indicates in the Copy State field whether the session has stopped.

10. **Suspend the B to C volume pairs.** This step stops the transfer of data between the B and C volume pairs.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> pausepprc -dev storage_image_ID_-remotedev storage_image_ID
SourceVolumeID:TargetVolumeID
```

Example

dscli> pausepprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 1200-125f:0700-075f

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

11. Wait until all of the out-of-sync (OOS) tracks have drained from the C and A volume pairs and the OOS count at C is zero. If you want to monitor this process, issue the lspprc command to query the status of the C to A volume pairs in Global Copy relationships.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID
 -l SourceVolumeID:TargetVolumeID

Example

```
dscli> lspprc -dev IBM.2107-1831860 -remotedev IBM.2107-75ALA2P -1 1200-125f:0700-075f
```

12. Suspend the C and A volume pairs.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> pausepprc -dev storage_image_ID_-remotedev storage_image_ID
SourceVolumeID:TargetVolumeID
```

Example

dscli> pausepprc -dev IBM.2107-1831760 -remotedev IBM.2107-130165X 1200-125f:0700-075f

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

13. End the Global Copy relationship between the B to C volumes at the C remote volume site.

Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -dev storage_image_ID -remotedev storage_image_ID
-at tgt -unconditional -quiet TargetVolumeID
Example

dscli> rmpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 -at tgt -unconditional -quiet 1A00-1A5f

14. **Reverse the direction by making the site A volumes a suspended primary site.** Use the **failoverpprc** command for A to C with cascading allowed and specifying Global Copy mode.

Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failoverpprc -dev storage_image_ID -remotedev storage_image_ID
 -type gcp -cascade SourceVolumeID:TargetVolumeID

Example

dscli> failoverpprc -dev IBM.2107-130165X -remotedev IBM.2107-1831760 -type gcp -cascade 1A00-1A5f:1200-125f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201 successfully reversed.

15. Resynchronize the A to C relationships.

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> failbackpprc -dev storage_image_ID -remotedev storage_image_ID
    -type gcp SourceVolumeID:TargetVolumeID
```

Example

dscli> failoverpprc -dev IBM.2107-130165X -remotedev IBM.2107-1831760
 -type gcp -cascade 1A00-1A5f:1200-125f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200 successfully failedback. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201 successfully failedback.

16. Establish Metro Mirror relationships between the A to B volumes using the incremental resynchronization function and the override option. As a result of this step, the relationship verification is bypassed and the incremental resynchronization function stops. The system determines which data to copy, so a full volume copy is bypassed and only changes are copied from the A to B Metro Mirror volume pairs.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -remotedev storage_image_ID -dev storage_image_ID
type mmir -mode nocp -incrementalresync override SourceVolumeID:TargetVolumeID

Example

dscli> mkpprc -remotedev IBM.2107-75ALA2P -dev IBM.2107-130165X -type mmir -mode nocp -incrementalresync override 2100-2107:2100-2107

The following represents an example of the output:

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.

17. Start incremental resynchronization with the initialization option on the B volumes in Metro Mirror relationships. Issue the mkpprc command at the intermediate site with the -incrementalresync enable parameter specified.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID
-type mmir -mode nocp -incrementalresync enable SourceVolumeID:TargetVolumeID

Example

dscli> mkpprc -dev IBM.2107-130165X -remotedev 2107-75ALA2P -type mmir -mode nocp -incrementalresync enable 2100-2107:2100-2107

The following represents an example of the output:

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.

18. Wait for the B to A volume pairs to reach the full duplex state. Issue the lspprc command if you want to monitor this activity.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID
-l SourceVolumeID:TargetVolumeID

Example

dscli> lspprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -1 1200-125f:1A00-1A5f

19. Start the Global Mirror session at the local site.

Enter the **mkgmir** command at the dscli command prompt with the following parameters and variables (from the local site):

dscli> mkgmir -dev storage_image_ID -lss LSS_ID -session session_ID

Example

dscli> mkgmir -remotedev IBM.2107-75ALA2P -lss 07 -session 31

The following represents an example of the output:

CMUC00162I mkgmir: Global Mirror for session 31 successfully started.

When this step is processed, the Metro/Global Mirror operations are running from site B to site A to site C. You are now ready to transition back to your original configuration, where site A is your production site.

20. Quiesce host I/O processing to the B volumes.

21. Suspend the B to A processing.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev storage_image_ID_-remotedev storage_image_ID
SourceVolumeID:TargetVolumeID

Example

dscli> pausepprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X 1200-125f:0700-075f

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

22. Create a relationship from the C volumes to the A volumes using the failoverpprc command with the -force and -cascade parameters specified.

Enter the failoverpprc command at the dscli prompt with the following parameters and variables:

dscli> failoverpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp -cascade -force SourceVolumeID:TargetVolumeID

Example

dscli> failoverpprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P -type gcp -cascade -force 1200-125f:1A00-1A5f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1200:1A00 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1201:1A01 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1202:1A02 successfully reversed.

23. End the Global Copy relationships between the B and A volume pairs at the local site.

Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -dev storage_image_ID -remotedev storage_image_ID -at tgt -unconditional -quiet TargetVolumeID

Example

```
dscli> rmpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P
-at tgt -unconditional -quiet 1A00-1A5f
```

The following represents an example of the output:

CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully withdrawn.

24. Resume host I/O processing to the A volumes.

25. Copy changes from site C back to site A, using the -force parameter.

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> failbackpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp -force SourceVolumeID:TargetVolumeID
```

Example

dscli> failoverpprc -dev IBM.2107-1831760 -remotedev IBM.2107-130165X -type gcp -force 1A00-1A5f:1200-125f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200 successfully failedback. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201 successfully failedback.

Note: Global Mirror processing continues to operate with site A volumes to site C volumes.

26. Wait for the first pass copy to complete from site C to site B. Issue the lspprc command if you want to monitor this activity and determine when the first pass status changes to "True."

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID
-l SourceVolumeID:TargetVolumeID

Example

dscli> lspprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P -l 1200-125f:1A00-1A5f

27. Modify Global Copy relationships between the A and C volume pairs. Specify the NOCOPY option and initiate incremental resynchronization without initialization.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp -incrementalresync enablenoinit -mode nocp
SourceVolumeID:TargetVolumeID
```

Example

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM2107-1831760 -type gcp -incrementalresync enablenoinit -mode nocp 1200-125f:1A00-1A5f

The following represents the first two lines of the output:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1200:1A00 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1201:1A01 successfully created.

28. Begin the process to include your B site in the 3-site Metro/Global Mirror configuration with production on site A. The Global Mirror session between the A, C, and D volumes must be stopped.

Enter the **rmgmir** command at the dscli command prompt with the following parameters and variables:

dscli> rmgmir -dev storage_image_ID -lss LSS_ID -session session_ID
Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID

Example

dscli> rmgmir -dev IBM.2107-130165X -quiet -lss 07 -session 2

The following represents an example of the output:

CMUC00165I pausegmir: Global Mirror for session 2 successfully paused.

29. Suspend the A to C volume pairs.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev storage_image_ID_-remotedev storage_image_ID SourceVolumeID:TargetVolumeID

Example

```
dscli> pausepprc -dev IBM.2107-130165X -remotedev IBM.2107-1831760 1200-125f:0700-075f
```

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

30. End the Global Copy relationships between the A to C volumes at the remote site.

Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -dev storage_image_ID -remotedev storage_image_ID -at tgt -unconditional -quiet TargetVolumeID

Example

```
dscli> rmpprc -dev IBM.2107-1831760 -remotedev IBM.2107-130165X
-at tgt -unconditional -quiet 1A00-1A5f
```

The following represents an example of the output:

CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully withdrawn.

31. Wait until all of the out-of-sync (OOS) tracks have drained from the C to B volume pairs and the OOS count is zero. If you want to monitor this process, issue the lspprc command to query the status of the C to B volume pairs in Global Copy relationships.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID
    -l SourceVolumeID:TargetVolumeID
```

Example

```
dscli> lspprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831860 -1 1200-125f:0700-075f
```

32. Suspend the C to B volume pairs.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev storage_image_ID_-remotedev storage_image_ID
SourceVolumeID:TargetVolumeID

Example

dscli> pausepprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P 1200-125f:0700-075f

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

33. Reverse the direction by making the site B volumes a suspended primary site. Use the

failoverpprc command for B to C specifying the Global Copy mode and that cascading is allowed. Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> failoverpprc -dev storage_image_ID -remotedev storage_image_ID
  -type gcp -cascade SourceVolumeID:TargetVolumeID
```

Example

dscli> failoverpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 -type gcp -cascade 1A00-1A5f:1200-125f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201 successfully reversed.

34. Resynchronize the C to B relationships.

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -remotedev storage_image_ID -dev storage_image_ID -type gcp SourceVolumeID:TargetVolumeID

Example

dscli> failoverpprc -remotedev IBM.2107-1831760 -dev IBM.2107-75ALA2P -type gcp -cascade 1A00-1A5f:1200-125f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200 successfully failedback. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201 successfully failedback.

35. Establish Metro Mirror relationships between the A to B volumes using the incremental resynchronization function and the override option. As a result of this step, the relationship verification is bypassed and the incremental resynchronization function stopped. The system determines which data to copy, so a full volume copy is bypassed and only changes are copied from the A to B Metro Mirror volume pairs.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID
type mmir -mode nocp -incrementalresync override SourceVolumeID:TargetVolumeID

Example

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type mmir -mode nocp -incrementalresync override 2100-2107:2100-2107

The following represents an example of the output:

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.

36. Start incremental resynchronization with the initialization option on the A volumes in the Metro Mirror relationships. Use the mkpprc command at the local site with the -incrementalresync enable parameter specified.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID
-type mmir -mode nocp -incrementalresync enable SourceVolumeID:TargetVolumeID

Example

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type mmir -mode nocp -incrementalresync enable 2100-2107:2100-2107

The following represents an example of the output:

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.

37. Wait for A to B to reach the full duplex state and for the first pass of the Global Copy processing of the B and C volumes to complete. You can monitor this activity by entering the lspprc command to query the status of the B to C volume pairs in Global Copy relationships.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-75ALA2P -l -fmt default 1200-125f:0700-075f

38. **Start Global Mirror at the intermediate site.** Now that the original infrastructure has been restored, you can resume the Global Mirror session.

Enter the **mkgmir** command at the dscli command prompt with the following parameters and variables:

dscli> mkgmir -dev IBM.2107-75ALA2P -session 1 -lss 07

The following represents an example of the output:

CMUC00164I resumegmir: Global Mirror for session 1 successfully resumed.

39. Verify that consistency groups are forming successfully.

Enter the **showgmir** -**metrics** command at the dscli command prompt with the following parameters and variables:

dscli> showgmir -metrics 07

The following represents an example of the output:

ID	Total Failed CG Count	Total Succes- sful CG Count	Succes- sful CG Percen- tage	Failed CG after Last Success	Last Succes- sful CG Form Time	Coord. Time (milli- seconds)	CG Interval Time (sec- onds)
IBM.2107 -130165X /07	0	55	100	0	10/20/ 2005 11:38:25 MST	50	0

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
(Seconds)	Unit	100	Status	Reason	State	Unit	100
30	-	-	No Error	-	-	-	-

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
No Error	-	-	-	-	No Error	-	-

Using forced failover and failback during an unplanned Metro/Global Mirror outage

Use this process to run failover operations (DS8000 only) from the local (A) site to the intermediate (B) site during an unplanned outage.

The following assumptions apply to your 3-site Metro/Global Mirror configuration:

- You used incremental resynchronization to establish the relationship between your site A and site B volumes and between your site A and site C volumes.
- You established a Global Mirror session at the local site. This means that Fibre Channel paths were established between all Global Mirror source and target pairs and between the master and subordinate storage units.

The command examples use the following site identifiers:

- Site A is identified as 2107-130165X
- Site B is identified as 2107-75ALA2P
- Site C is identified as 2107-183176O

The process described in this task provides a disaster recovery solution when an unplanned failure occurs at your local site and you want to limit the amount of interruption to your local production processing. You can run your operations from your intermediate site, which is protected by a two-site Global Mirror configuration, until your local site recovers. Global Mirror continues sending updates to the storage unit at the remote site and continues to form consistency groups.

This process uses forced failover and forced failback processing to establish a relationship between the C and A volumes without verification that this relationship already existed. The **failoverpprc** and **failbackpprc** commands with the **-force** parameter accommodate this processing.

Attention: Use the **-force** parameter only as directed. This parameter can cause severe damage to your data if it is misused. Contact IBM Support before you attempt to use the **-force** parameter if your outage situation is outside the boundaries of this example.

Complete the following steps to failover and restore operations to the intermediate (B) site.

1. At the local site, ensure that data consistency is achieved between the site A and site B volumes. If the local site was not completely destroyed, it is essential that data from any surviving A and B volume pairs be copied and a consistent copy be achieved at the remote site. You can use freeze and unfreeze commands that are supported using external automation software to create data consistency to multiple Metro Mirror volume pairs.

To freeze write activity to Metro Mirror primary volumes, complete the following steps:

a. Freeze updates to the A volumes in Metro Mirror relationships across the affected LSSs. This ensures that the B volumes will be consistent at the time of the freeze process. (One command per LSS is required.)

Enter the **freezepprc** command at the dscli command prompt with the following parameters and variables:

dscli> freezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07-12

The following represents an example of the output:

CMUC00161W freezepprc: Remote Mirror and Copy consistency group 07:12 successfully created.

As a result of the freeze action, the following processing occurs:

- The established Remote Mirror and Copy paths between the LSS pairs are deleted.
- The volume pairs that are associated with the source and target LSSs are suspended. During this time, the storage unit collects data that is sent to the A volumes in Metro Mirror relationships.
- I/O to the Metro Mirror volume pairs is temporarily queued.
- b. **Resume operations following a freeze.** This operation—also called a thaw operation—allows I/O processing to resume for the specified volume pairs.

Enter the **unfreezepprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> unfreezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P
07:12
```

The following represents an example of the output:

 $\ensuremath{\mathsf{CMUC00198I}}$ unfreezepprc: Remote Mirror and Copy pair 07:12 successfully thawed.

2. Create a relationship from the C volumes to the A volumes, using the -force and -cascade parameters. No validation is done at site C to determine that site C is a secondary of site A.

Note: For this step to succeed you must ensure that the Remote Mirror and Copy paths between all Global Mirror source and target pairs and between the Master and subordinate storage units have been created.

Enter the **failoverpprc** command at the dscli prompt with the following parameters and variables:

dscli> failoverpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp -cascade -force SourceVolumeID:TargetVolumeID

Example

dscli> failoverpprc -dev IBM.2107-1831760 -remotedev IBM.2107-130165X -type gcp -cascade -force 1200-125f:1A00-1A5f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1200:1A00 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1201:1A01 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1202:1A02 successfully reversed.

3. End the Metro Mirror relationship between the A to B volumes at the B volumes intermediate site.

Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -dev storage_image_ID -remotedev storage_image_ID -at tgt -unconditional -quiet TargetVolumeID

Example

dscli> rmpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -at tgt -unconditional -quiet 1200-125f

4. **Redirect host I/O processing to the B volumes.** Changes are recorded on the B volumes until the A volumes can be resynchronized with the B volumes. Also, Global Mirror continues to operate from site B to site C.

Note: You can run in this configuration until the A site has recovered and you want to restore operations there. Begin the next step after the A volumes have been recovered and you're still in production on the B volumes.

5. **Copy changes from site C back to site A, using the -force parameter.** Host I/O processing continues uninterrupted to the B volumes while the A volumes are made current. (The data is still flowing from B to C, so any changes made to B are being transferred to C and therefore will get from C to A.) This command copies the changes back to the A volumes that were made to the B volumes while hosts were running on the A volumes. (In a DS CLI environment, where the local and remote sites are not using the same management console, you have to use the management console of the remote site.)

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp -force SourceVolumeID:TargetVolumeID

Example

dscli> failbackpprc -dev IBM.2107-1831760 -remotedev IBM.2107-130165X -type gcp -force 1200-125f:1A00-1A5f

The following represents an example of the output:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1200:1A00 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1201:1A01 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1202:1A02 successfully failed back.

6. Wait for the first pass copy to complete from site C to site A. Issue the lspprc command if you want to monitor this activity and determine when the first pass status changes to "True."

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID
-l SourceVolumeID:TargetVolumeID
```

Example

```
dscli> lspprc -dev IBM.2107-1831760 -remotedev IBM.2107-75130165X -l 1200-125f:1A00-1A5f
```

The following represents the first two lines of the report generated by the **lspprc** command:

ID	State	Reason	Туре	Source- LSS	Time- out (secs)	Critical Mode	First Pass Status
IBM.2107- 183176O /2101: IBM.2107- 130165X /2101	Copy Pend- ing	-	Global Copy	IBM.2107- 130165X /20	300	Disabled	True
IBM.2107- 183176O /2100: IBM.2107- 130165X /2100	Copy Pend- ing	-	Global Copy	IBM.2107- 130165X /20	300	Disabled	True

7. **Modify Global Copy relationships between the B and C volume pairs.** Specify the NOCOPY option and initiate incremental resynchronization without initialization.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp -incrementalresync enablenoinit -mode nocp
SourceVolumeID:TargetVolumeID
```

Example

dscli> mkpprc -dev IBM.2107-75ALA2P -remotedev IBM2107-1831760 -type gcp -incrementalresync enablenoinit -mode nocp 1200-125f:1A00-1A5f

The following represents the first two lines of the output:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1200:1A00 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1201:1A01 successfully created.

8. Begin the process to return production to site A. First, the Global Mirror session at site B must be stopped.

Enter the **rmgmir** command at the dscli command prompt with the following parameters and variables:

dscli> rmgmir -dev storage_image_ID -lss LSS_ID -session session_ID
Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID

Example

dscli> rmgmir -dev IBM.2107-75ALA2P -quiet -lss 07 -session 1

The following represents an example of the output:

CMUC00165I rmgmir: Global Mirror for session 1 successfully terminated.

9. Verify that the Global Mirror session has ended.

Enter the **showgmir** command at the dscli command prompt with the following parameters and variables:

dscli> showgmir -dev storage_image_ID LSS_ID

Example

dscli> showgmir -dev IBM.2107-75ALA2P 10

In the resulting report, the output indicates in the Copy State field whether the session has stopped.

10. **Suspend the B to C volume pairs.** This step stops the transfer of data between the B and C volume pairs.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev storage_image_ID_-remotedev storage_image_ID SourceVolumeID:TargetVolumeID

Example

dscli> pausepprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 1200-125f:0700-075f

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

11. Wait until all of the out-of-sync (OOS) tracks have drained from the C and A volume pairs and the OOS count at C is zero. If you want to monitor this process, issue the lspprc command to query the status of the C to A volume pairs in Global Copy relationships.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID
 -1 SourceVolumeID:TargetVolumeID

Example

dscli> lspprc -dev IBM.2107-1831860 -remotedev IBM.2107-75ALA2P -1 1200-125f:0700-075f

12. Suspend the C and A volume pairs.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev storage_image_ID_-remotedev storage_image_ID
SourceVolumeID:TargetVolumeID

Example

dscli> pausepprc -dev IBM.2107-1831760 -remotedev IBM.2107-130165X 1200-125f:0700-075f

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

13. End the Global Copy relationship between the B to C volumes at the C remote volume site.

Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -dev storage_image_ID -remotedev storage_image_ID
-at tgt -unconditional -quiet TargetVolumeID

Example

dscli> rmpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 -at tgt -unconditional -quiet 1A00-1A5f

14. **Reverse the direction by making the site A volumes a suspended primary site.** Use the **failoverpprc** command for A to C with cascading allowed and specifying Global Copy mode.

Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failoverpprc -dev storage_image_ID -remotedev storage_image_ID
 -type gcp -cascade SourceVolumeID:TargetVolumeID

Example

```
dscli> failoverpprc -dev IBM.2107-130165X -remotedev IBM.2107-1831760
  -type gcp -cascade 1A00-1A5f:1200-125f
```

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201 successfully reversed.

15. Resynchronize the A to C relationships.

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp SourceVolumeID:TargetVolumeID

Example

dscli> failoverpprc -dev IBM.2107-130165X -remotedev IBM.2107-1831760 -type gcp -cascade 1A00-1A5f:1200-125f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200 successfully failedback. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201 successfully failedback.

16. Establish Metro Mirror relationships between the A to B volumes using the incremental resynchronization function and the override option. As a result of this step, the relationship verification is bypassed and the incremental resynchronization function stops. The system determines which data to copy, so a full volume copy is bypassed and only changes are copied from the A to B Metro Mirror volume pairs.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -remotedev storage_image_ID -dev storage_image_ID
type mmir -mode nocp -incrementalresync override SourceVolumeID:TargetVolumeID

Example

dscli> mkpprc -remotedev IBM.2107-75ALA2P -dev IBM.2107-130165X -type mmir -mode nocp -incrementalresync override 2100-2107:2100-2107

The following represents an example of the output:

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.

17. Start incremental resynchronization with the initialization option on the B volumes in Metro Mirror relationships. Issue the mkpprc command at the intermediate site with the -incrementalresync enable parameter specified.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID
-type mmir -mode nocp -incrementalresync enable SourceVolumeID:TargetVolumeID

Example

dscli> mkpprc -dev IBM.2107-130165X -remotedev 2107-75ALA2P -type mmir -mode nocp -incrementalresync enable 2100-2107:2100-2107

The following represents an example of the output:

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.

18. Wait for the B to A volume pairs to reach the full duplex state. Issue the lspprc command if you want to monitor this activity.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID
-l SourceVolumeID:TargetVolumeID

Example

dscli> lspprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -l 1200-125f:1A00-1A5f

19. Start the Global Mirror session at the local site.

Enter the **mkgmir** command at the dscli command prompt with the following parameters and variables (from the local site):

dscli> mkgmir -dev storage_image_ID -lss LSS_ID -session session_ID

Example

dscli> mkgmir -remotedev IBM.2107-75ALA2P -lss 07 -session 31

The following represents an example of the output:

CMUC00162I mkgmir: Global Mirror for session 31 successfully started.

When this step is processed, the Metro/Global Mirror operations are running from site B to site A to site C. You are now ready to transition back to your original configuration, where site A is your production site.

20. Quiesce host I/O processing to the B volumes.

21. Suspend the B to A processing.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev storage_image_ID_-remotedev storage_image_ID
SourceVolumeID:TargetVolumeID

Example

```
dscli> pausepprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X 1200-125f:0700-075f
```

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

22. Create a relationship from the C volumes to the A volumes using the failoverpprc command with the -force and -cascade parameters specified.

Enter the **failoverpprc** command at the dscli prompt with the following parameters and variables:

dscli> failoverpprc -dev storage_image_ID -remotedev storage_image_ID -type gcp -cascade -force SourceVolumeID:TargetVolumeID

Example

```
dscli> failoverpprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P -type gcp -cascade -force 1200-125f:1A00-1A5f
```

The following represents an example of the output:

```
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1200:1A00
successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1201:1A01
successfully reversed.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1202:1A02
successfully reversed.
```

23. End the Global Copy relationships between the B and A volume pairs at the local site.

Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -dev storage_image_ID -remotedev storage_image_ID
-at tgt -unconditional -quiet TargetVolumeID

Example

```
dscli> rmpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P
-at tgt -unconditional -quiet 1A00-1A5f
```

The following represents an example of the output:

CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully withdrawn.

24. Resume host I/O processing to the A volumes.

25. Copy changes from site C back to site A, using the -force parameter.

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> failbackpprc -dev storage_image_ID -remotedev storage_image_ID
  -type gcp -force SourceVolumeID:TargetVolumeID
```

Example

dscli> failoverpprc -dev IBM.2107-1831760 -remotedev IBM.2107-130165X -type gcp -force 1A00-1A5f:1200-125f

The following represents an example of the output:

```
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200
successfully failedback.
CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201
successfully failedback.
```

Note: Global Mirror processing continues to operate with site A volumes to site C volumes.

26. Wait for the first pass copy to complete from site C to site B. Issue the **lspprc** command if you want to monitor this activity and determine when the first pass status changes to "True."

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID
-l SourceVolumeID:TargetVolumeID
```

Example

dscli> lspprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P -l 1200-125f:1A00-1A5f

27. Modify Global Copy relationships between the A and C volume pairs. Specify the NOCOPY option and initiate incremental resynchronization without initialization.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp -incrementalresync enablenoinit -mode nocp
SourceVolumeID:TargetVolumeID
```

Example

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM2107-1831760 -type gcp -incrementalresync enablenoinit -mode nocp 1200-125f:1A00-1A5f

The following represents the first two lines of the output:

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1200:1A00 successfully created. CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 1201:1A01 successfully created.

28. Begin the process to include your B site in the 3-site Metro/Global Mirror configuration with production on site A. The Global Mirror session between the A, C, and D volumes must be stopped.

Enter the **rmgmir** command at the dscli command prompt with the following parameters and variables:

dscli> rmgmir -dev storage_image_ID -lss LSS_ID -session session_ID
Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID

Example

dscli> rmgmir -dev IBM.2107-130165X -quiet -lss 07 -session 2

The following represents an example of the output:

CMUC00165I pausegmir: Global Mirror for session 2 successfully paused.

29. Suspend the A to C volume pairs.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev storage_image_ID_-remotedev storage_image_ID
SourceVolumeID:TargetVolumeID

Example

dscli> pausepprc -dev IBM.2107-130165X -remotedev IBM.2107-1831760 1200-125f:0700-075f

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

30. End the Global Copy relationships between the A to C volumes at the remote site.

Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> rmpprc -dev storage_image_ID -remotedev storage_image_ID
-at tgt -unconditional -quiet TargetVolumeID
```

Example

dscli> rmpprc -dev IBM.2107-1831760 -remotedev IBM.2107-130165X -at tgt -unconditional -quiet 1A00-1A5f

The following represents an example of the output:

CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully withdrawn.

31. Wait until all of the out-of-sync (OOS) tracks have drained from the C to B volume pairs and the OOS count is zero. If you want to monitor this process, issue the lspprc command to query the status of the C to B volume pairs in Global Copy relationships.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev storage_image_ID -remotedev storage_image_ID
 -l SourceVolumeID:TargetVolumeID

Example

dscli> lspprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831860 -1 1200-125f:0700-075f

32. Suspend the C to B volume pairs.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev storage_image_ID_-remotedev storage_image_ID SourceVolumeID:TargetVolumeID

Example

dscli> pausepprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P 1200-125f:0700-075f

The following represents an example of the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1200:0700 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 1201:0701 relationship successfully paused.

33. Reverse the direction by making the site B volumes a suspended primary site. Use the **failoverpprc** command for B to C specifying the Global Copy mode and that cascading is allowed.

Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failoverpprc -dev storage_image_ID -remotedev storage_image_ID
 -type gcp -cascade SourceVolumeID:TargetVolumeID

Example

dscli> failoverpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760
 -type gcp -cascade 1A00-1A5f:1200-125f

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201 successfully reversed.

34. Resynchronize the C to B relationships.

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

```
dscli> failbackpprc -remotedev storage_image_ID -dev storage_image_ID
-type gcp SourceVolumeID:TargetVolumeID
```

Example

```
dscli> failoverpprc -remotedev IBM.2107-1831760 -dev IBM.2107-75ALA2P
-type gcp -cascade 1A00-1A5f:1200-125f
```

The following represents an example of the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A00:1200 successfully failedback. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1A01:1201 successfully failedback.

35. Establish Metro Mirror relationships between the A to B volumes using the incremental resynchronization function and the override option. As a result of this step, the relationship verification is bypassed and the incremental resynchronization function stopped. The system determines which data to copy, so a full volume copy is bypassed and only changes are copied from the A to B Metro Mirror volume pairs.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID
type mmir -mode nocp -incrementalresync override SourceVolumeID:TargetVolumeID

Example

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type mmir -mode nocp -incrementalresync override 2100-2107:2100-2107

The following represents an example of the output:

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.

36. Start incremental resynchronization with the initialization option on the A volumes in the Metro Mirror relationships. Use the mkpprc command at the local site with the -incrementalresync enable parameter specified.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID
-type mmir -mode nocp -incrementalresync enable SourceVolumeID:TargetVolumeID

Example

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type mmir -mode nocp -incrementalresync enable 2100-2107:2100-2107

The following represents an example of the output:

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.

37. Wait for A to B to reach the full duplex state and for the first pass of the Global Copy processing of the B and C volumes to complete. You can monitor this activity by entering the lspprc command to query the status of the B to C volume pairs in Global Copy relationships.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-75ALA2P -l -fmt default 1200-125f:0700-075f

38. **Start Global Mirror at the intermediate site.** Now that the original infrastructure has been restored, you can resume the Global Mirror session.

Enter the **mkgmir** command at the dscli command prompt with the following parameters and variables:

dscli> mkgmir -dev IBM.2107-75ALA2P -session 1 -lss 07

The following represents an example of the output:

CMUC00164I resumegmir: Global Mirror for session 1 successfully resumed.

39. Verify that consistency groups are forming successfully.

Enter the **showgmir** -**metrics** command at the dscli command prompt with the following parameters and variables:

dscli> showgmir -metrics 07

The following represents an example of the output:

ID	Total Failed CG Count	Total Succes- sful CG Count	Succes- sful CG Percen- tage	Failed CG after Last Success	Last Succes- sful CG Form Time	Coord. Time (milli- seconds)	CG Interval Time (sec- onds)
IBM.2107 -130165X /07	0	55	100	0	10/20/ 2005 11:38:25 MST	50	0

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
30	-	-	No Error	-	-	-	-

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
No Error	-	-	-	-	No Error	-	-

Discarding changes or committing changes to consistency groups

Use this process to determine whether to discard or commit changes to FlashCopy volumes (DS8000 only) that are part of consistency groups.

When you query the state of the consistency group, the output displays whether the sequence numbers are equal for all FlashCopy relationships that are part of the consistency groups and whether the FlashCopy relationships are revertible.

If the sequence numbers are not equal, which results from something going wrong during a FlashCopy consistency group formation operation, you must determine the action to take. The action that you take depends on which phase the consistency group formation was in at the time of the failure. For example, if the failure occurred while FlashCopy commands were processing, intervention is required to provide the consistency. The action depends on the current status of the FlashCopy, where the sequence numbers and the revertible state are important.

The following options are available to help you determine the action to take:

• **Discard changes (revert to a previous consistent state).** Assume that the sequence numbers of the FlashCopy relationships are different and the copy process has not started for all the volumes. In this case, the FlashCopy data is inconsistent and cannot be used. You must revert changes, which removes all not-committed data from the FlashCopy target and reverts (or is restored) to the last consistency group.

Note: You can discard changes to FlashCopy target volumes *only* if you have modified the FlashCopy relationship using the **setflashrevertible** command, which changes the Revertible value to Enabled. When you revert a FlashCopy relationship that is in a revertible state, ensure that you specify its associated FlashCopy sequence number.

• Commit all FlashCopy relationships in the consistency group to the current level.

Assume that the sequence numbers are all equal and there is a mix of revertible and nonrevertible volumes and the copy process to the FlashCopy target volumes has occurred but not completed for some volumes. In this case, the FlashCopy target volumes are usable and the process has to be committed manually.

This is done by issuing a commit command to *all* revertible FlashCopy relationships to commit data to the FlashCopy target volumes and create data consistency between the source and target volumes. The commit process specifies that the last consistency group that has been created by the Global Mirror session is committed to the current state, and reverting to the previous consistency group state is no longer possible.

Note: You can commit changes to FlashCopy target volumes *only* if you have modified the FlashCopy relationship using the **setflashrevertible** command, which changes the Revertible value to Enabled.

Recovery scenario using incremental resynchronization in a Metro/Global Mirror configuration

Use this process (DS8000 only) to restart recovery using the incremental resynchronization function during an outage at the intermediate site.

In a Metro/Global Mirror configuration, if you lose access to the storage unit at the intermediate site (either in a planned or unplanned outage), you can restart a two-site Global Mirror environment between the local and remote sites. You can use the incremental resynchronization function to avoid having to run a full copy of the volumes from the local site to the remote site.

This scenario describes the steps for restarting the recovery environment running Global Mirror from the local site to the remote site using the incremental resynchronization function. For best management

practices, combine the functions of a Metro/Global Mirror environment with automation such as Geographically Dispersed Parallel Sysplex (GDPS) to ensure continuous or near-continuous availability during outages, including disasters.

Notes:

- The following assumptions are made before you initiate the steps in this scenario:
 - You have established all your Remote Mirror and Copy paths before you establish your pairs or initiate any of the incremental resynchronization process. If the paths are not established first, an error condition might result.
 - You have established your Metro Mirror volume pairs to use the incremental resynchronization function on each of the primary volumes when you configured your Metro/Global Mirror configuration.
 - You have specified the -mode full parameter for each of these volume pairs.

Notes:

- The command parameters and options that are used in this scenario are examples.
- Some of the query output is presented in table format for clarity. The actual report is not displayed in this format.
- The output for some commands differs depending on the storage unit from which you issue the command.

Complete these steps for the recovery operation:

1. Enable the incremental resynchronization option for the A to B Metro Mirror volume pairs. If this is the first attempt to establish the volume pairs, specify **-mode full** as shown in the **mkpprc** command example. Otherwise, specify **-mode nocp**.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type mmir -mode full -incrementalresync enable 2100-2107:2100-2107

The following example represents the output:

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.

See "Creating a Metro Mirror relationship" on page 569 for more information.

2. Pause (suspend) all A to B Metro Mirror volume pairs. Some (but not all) volume pairs might have been suspended with the outage of the intermediate site.

Note: If the consistency group function is being used, the automation application (such as GPDS) issued the **freezepprc** command and all devices are suspended.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -unconditional -at src 2100-2107

The following example represents the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 2100 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 2101 relationship successfully paused.

See "Pausing a Metro Mirror relationship" on page 572 for more information.

Notes:

- a. With the volume pairs suspended, updates to the A volumes are marked in the change recording and out-of-synchronization bitmaps on the Metro Mirror A volumes at the local site.
- b. The master storage unit might have been in the process of using FlashCopy to copy the consistency group to the D volumes when the outage occurred and the consistency group formation was not able to complete. If so, you must verify the consistency group formation. See "Querying Global Mirror processing" on page 588 for more information.
- **3. Issue a failover command to the C to B volumes at the remote site, specifying the -cascade option:** With the loss of the B volumes at the intermediate site, the state of the C volumes is changed from secondary duplex pending (or suspended) to Suspended Host Source when the command processes. Updates are collected in out-of-sync bitmaps.

Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failoverpprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P -type gcp -cascade 2100-2107:2100-2107

The following example represents the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 2100:2100 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 2101:2101 successfully reversed.

See "Running a recovery failover operation" on page 578 for more information.

4. After the failover operation, you can view the status of the volumes to determine the state of the volumes: From the remote site, enter the lspprc command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -l 2100-2107

The following example represents the output:

ID	State	Reason	Туре	Out of Sync Tracks	Tgt Read	Src Cascade
2100:2100	Suspended Host Source	-	Global Copy	0	Disabled	Enabled
2101:2101	Suspended Host Source	-	Global Copy	0	Disabled	Enabled

Tgt Cas cade	Date Sus- pended	Source LSS	Time- out (secs)	Crit Mode	First Pass Status	Incre- mental Resync	Tgt Write
Invalid	-	21	300	Disabled	True	Disabled	Disabled
Invalid	-	21	300	Disabled	True	Disabled	Disabled

5. Attempt to clean up any surviving components of Global Mirror at the intermediate site, if needed.

a. End the Global Mirror session at the master storage unit.

Enter the **rmgmir** command at the dscli command prompt with the following parameters and variables (from the intermediate site):

dscli> rmgmir -dev IBM.2107-75ALA2P -quiet -lss 20 -session 31

The following example represents the output:

 $\ensuremath{\mathsf{CMUC00165I}}$ rmgmir: Global Mirror for session 31 successfully stopped.

- b. End the Global Mirror session at the subordinate storage units. Reissue the command if the Global Mirror session does not stop because of subordinate storage units still associated to the master storage unit. See "Ending a Global Mirror session" on page 597 for more information.
- 6. Verify the Global Mirror consistency group formation: If the intermediate site outage occurred in the middle of consistency group formation, you must determine whether the FlashCopy operations must be committed or reverted.

Enter the **lsflash** command at the dscli command prompt with the following parameters and variables.

dscli> lsflash -l 2100-2107

See "Viewing information about FlashCopy relationships" on page 556 for more information. The following table represents an example of the output:

ID	SrcLSS	Seq- uence Num	Timeout	Active Copy	Recording	Persistent
2100:2300	21	44357D55	300	Disabled	Enabled	Enabled
2101:2301	21	44357D55	300	Disabled	Enabled	Enabled

Revertible	Source- Write Enabled	Target- Write Enabled	Back- ground Copy	Out Of Sync Tracks	Date Created	Date- Synced
Disabled	Enabled	Disabled	Disabled	1525879	Fri Mar 24 09:45:54 MST 2006	Thu Apr 06 13:42:58 MST 2006
Disabled	Enabled	Disabled	Disabled	1525879	Fri Mar 24 09:45:54 MST 2006	Thu Apr 06 13:42:58 MST 2006

7. Establish Global Copy relationships using the A and C volume pairs with the Incremental **Resynchronization recover option:** Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

See "Creating a Global Copy relationship" on page 572 for more information.

The following example represents the output:

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.

Notes:

- a. The C volumes were primary suspended volumes that had Global Copy relationships with the B volumes, which were in Metro Mirror relationships with the A volumes.
- b. The Incremental Resynchronization function that is running on the A volumes is stopped. The tracks of data in the change recording and out-of-synchronization bitmaps are merged and copied from the A volumes to the C volumes.
- 8. Wait for the first pass of Global Copy processing to complete between the A to C volumes: You can monitor this activity by querying the status of the volumes.

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM.2107-1831760 -type gcp -incrementalresync recover 2100-2107:2100-2107

From the local site, enter the **lspprc** command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev IBM.2107-130165X -remotedev IBM.2107-1831760 2101:2101 2100:2101

The following example represents the output:

ID	State	Reason	Туре	Source- LSS	Time- out (secs)	Critical Mode	First Pass Status
IBM.2107- 130165X /2101: IBM.2107- 1831760 /2101	Copy Pend- ing	-	Global Copy	IBM.2107- 1831760 /20	300	Disabled	True
IBM.2107- 130165X /2100: IBM.2107- 1831760 /2100	Copy Pend- ing	-	Global Copy	IBM.2107- 1831760 /20	300	Disabled	True

9. When the first pass of Global Copy processing is completed, start the Global Mirror session on the A volumes.

The master storage unit begins forming consistency groups for the specified Global Mirror session. Global Mirror runs from the local site to the remote site until the intermediate site is ready to resume operation.

Enter the **mkgmir** command at the dscli command prompt with the following parameters and variables (from the local site):

mkgmir -dev IBM.2107-130165X -lss 07 -session 31

The following example represents the output:

CMUC00162I mkgmir: Global Mirror for session 31 successfully started.

See "Starting Global Mirror processing" on page 590 for more information. When the intermediate site has been recovered, the volumes at the intermediate site must be resynchronized with the local volumes.

During the outage, data was written to the volumes at the local site. After the intermediate site is recovered, the volumes at the intermediate site must be resynchronized.

The former Metro/Global Mirror configuration must be "cleaned up" to reestablish it back to its original configuration. A host connection to the storage unit at the intermediate site is required.

10. Complete the following steps in preparation for a failback operation from the remote site to the intermediate site:

a. End the Metro Mirror relationship between the A to B volumes at the intermediate site. Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -quiet -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -unconditional -at tgt 2100-2107

The following example represents the output:

CMUC00155I rmpprc: Remote Mirror and Copy volume pair :2100 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair :2101 relationship successfully withdrawn.

See "Deleting a Metro Mirror relationship" on page 573 for more information.

b. Pause (suspend) the B to C volume pairs if they are not already suspended. You can query the status of the volumes for this determination. From the remote site, enter the lspprc command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev IBM.2107-130165X -remotedev IBM.2107-1831760 2101:2101 2100:2101

The following example represents the output: See "Pausing a Metro Mirror relationship" on page 572 for more information.

ID	State	Reason	Туре	Source- LSS	Time- out (secs)	Critical Mode	First Pass Status
IBM.2107- 1831760 /2100: IBM.2107- 75ALA2P /2100	Suspend- ed Host Source	-	Global Copy	IBM.2107- 75ALA2P /21	unknown	Disabled	True
IBM.2107- 1831760 2101: IBM.2107- 75ALA2P /2101	Copy Pend- ing	-	Global Copy	IBM.2107- 1831760 /21	300	Disabled	True

If necessary, clean up the former Global Mirror configuration at the intermediate site using the following two steps:

c. End the Global Mirror session from the master storage unit at the intermediate site.

Note: If the Global Mirror session was successfully stopped at the time of the outage, this step might not be necessary and it might generate an error message when the command processes.

Enter the **rmgmir** command at the dscli command prompt with the following parameters and variables (from the intermediate site):

dscli> rmgmir -dev IBM.2107-1301261 -quiet -lss 20 -session 31

The following example represents the output:

CMUC00165I rmgmir: Global Mirror for session 31 successfully stopped.

See "Ending Global Mirror processing (script mode)" on page 590 or "Ending Global Mirror processing (no script)" on page 591 for more information.

d. If required, stop the Global Mirror session that is running from any of the subordinates.

Enter the **rmgmir** command at the dscli command prompt with the following parameters and variables:

dscli> rmgmir -quiet -lss 20 -session 31

The following example represents the output:

CMUC00165I rmgmir: Global Mirror for session 31 successfully stopped.

See "Ending Global Mirror processing (script mode)" on page 590 or "Ending Global Mirror processing (no script)" on page 591 for more information.

11. From the remote site, a failback Global Copy operation between the C to B volumes:

When the **failbackpprc** command processes, data will be copied from the remote site to the intermediate site. Specify the C volumes as the sources and the B volumes as targets with the failback command.

Note: Ensure the availability of the paths from the remote site to the intermediate site with the **lspprcpath** command.

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P -type gcp -cascade 2100-2107:2100-2107

The following example represents the output:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A00:1200 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1A01:1201 successfully failed back.

See "Running a recovery failback operation" on page 577 for more information.

12. Wait for the first pass to complete between the C volumes at the remote site and the B volumes at the intermediate site: You can monitor this activity by querying the status of the volumes.

Enter the **lspprc** command at the dscli command prompt with the following parameters and variables: (from the intermediate site)

```
dscli> lspprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P -l -fullid
-fmt default 2100-2107
```

See "Querying Global Mirror processing" on page 588 for more information.

ID	State	Reason	Туре	Out of Sync Tracks	Tgt Read	Src Cascade
IBM.2107- 1831760 /2100: IBM.2107- 75ALA2P /2100	Copy Pending	-	Global Copy	0	Disabled	Enabled
IBM.2107- 1831760 /2101: IBM.2107- 75ALA2P /2101	Copy Pending	-	Global Copy	0	Disabled	Enabled

The following example represents the output:

Tgt Cascade	Date Sus- pended	Source LSS	Time- out (secs)	Crit Mode	First Pass Status	Incre- mental Resync	Tgt Write
Invalid	-	IBM.2107- 1831760 /21	Unknown	Disabled	True	Disabled	Disabled
Invalid	-	IBM.2107- 1831760 /21	Unknown	Disabled	True	Disabled	Disabled

13. Start the Incremental Resynchronization function *without* the initialization option on the A volumes: This step allows you to "force" a resynchronization later between primary (A) volumes at the local site and the volumes at the intermediate site to ensure all updates are copied.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type gcp -incrementalresync enablenoinit -mode nocp 2100-2107:2100-2107

The following example represents the output:

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.

See "Creating a Metro Mirror relationship" on page 569 for more information. You are now ready to restore the original configuration Metro/Global Mirror without interrupting production.

14. Stop the Global Mirror session between the A and C volumes between the local and remote sites. During this transition time, the data on the D volumes in FlashCopy relationships might be consistent but not current until the transition is complete.

Enter the **rmgmir** command at the dscli command prompt with the following parameters and variables:

dscli> rmgmir -dev IBM.2107-130165X -quiet -lss 21 -session 31

The following example represents the output:

CMUC00165I rmgmir: Global Mirror for session 31 successfully stopped.

See "Ending Global Mirror processing (script mode)" on page 590 or "Ending Global Mirror processing (no script)" on page 591 for more information.

15. Allow the resynchronization of the C to B volumes to run by completing the following steps:

a. Pause (suspend) the A to C volume pairs that were established in Global Copy mode.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev IBM.2107-130165X -remotedev IBM.2107-1831760 2100-2107:2100-2107

The following example represents the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 2100:2100 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 2101:2101 relationship successfully paused.

See "Pausing a Metro Mirror relationship" on page 572 for more information.

b. Wait for data to be copied from the C volumes at the remote site to the B volumes at the intermediate site. Enter the lspprc command at the dscli command prompt with the following parameters and variables:

dscli> lspprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P -1 -fmt default 2100-2107

See "Querying Global Mirror processing" on page 588 for more information.

The following example represents the output:

ID	State	Reason	Туре	Out of Sync Tracks	Tgt Read	Src Cascade
IBM.2107- 1831760 /2100: IBM.2107- 75ALA2P /2100	Copy Pending	-	Global Copy	0	Disabled	Enabled

ID	State	Reason	Туре	Out of Sync Tracks	Tgt Read	Src Cascade
IBM.2107- 1831760 /2101: IBM.2107- 75ALA2P /2101	Copy Pending	-	Global Copy	0	Disabled	Enabled

Tgt Cascade	Date Sus- pended	Source LSS	Time- out (secs)	Crit Mode	First Pass Status	Incre- mental Resync	Tgt Write
Invalid	-	IBM.2107- 1831760 /21	Unknown	Disabled	True	Disabled	Disabled
Invalid	-	IBM.2107- 1831760 /21	Unknown	Disabled	True	Disabled	Disabled

c. End the A and C Global Copy relationship at the remote site.

Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

dscli> rmpprc -quiet -dev IBM.2107-1831760 -unconditional -at tgt 2100-2107

The following example represents the output:

CMUC00155I rmpprc: Remote Mirror and Copy volume pair :2100 relationship successfully withdrawn. CMUC00155I rmpprc: Remote Mirror and Copy volume pair :2101 relationship successfully withdrawn.

See "Removing the Global Copy pair relationship" on page 599 for more information.

Notes:

- 1) The value for the -dev parameter must be the remote site server (site C).
- 2) The management console must be able to communicate with the remote server for this command to process successfully.

When the command processes, the C volumes at the remote site are no longer the secondary volumes in a Global Copy relationship with the A volumes. This process allows for a later failback operation for the B to C volume pairs.

The Global Copy relationship between the A to C volumes was stopped at the remote site, which did not affect the status of the A volumes at the local site. The updates on the A volumes continue until the volumes are again fully synchronized.

16. After data on the C volumes has been copied to the B volumes, pause (suspend) the C to B volume pairs. This step is required before a failback operation can be issued between the B to C volumes, which requires the C volumes to be paused.

Enter the **pausepprc** command at the dscli command prompt with the following parameters and variables:

dscli> pausepprc -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P 2100-2107:2100-2107

The following example represents the output:

CMUC00157I pausepprc: Remote Mirror and Copy volume pair 2100:2100 relationship successfully paused. CMUC00157I pausepprc: Remote Mirror and Copy volume pair 2101:2101 relationship successfully paused.

See "Pausing a Metro Mirror relationship" on page 572 for more information.

17. At the intermediate site, issue a failover Global Copy operation to the B to C volumes, with the -cascade option: The B volumes are primary suspended volumes.

Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failoverpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 -type gcp -cascade 2100-2107:2100-2107

The following example represents the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 2100:2100 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 2101:2101 successfully reversed.

See "Running a recovery failover operation" on page 578 for more information.

18. At the intermediate site, run a failback Global Copy operation for the B to C volumes, with the -cascade option:

Enter the **failbackpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failbackpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831760 -type gcp -cascade 2100-2107:2100-2107

The following example represents the output:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 2100:2100 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 2101:2101 successfully failed back.

See "Running a recovery failback operation" on page 577 for more information.

19. Establish Metro Mirror relationships between the A to B volumes using the incremental resynchronization function and the override option. As a result, the relationship verification is bypassed and the incremental resynchronization function stopped. The change recording and out-of-synchronization bitmaps that were monitored and tracked on the primary Metro Mirror volumes are merged to determine the data to copy from the A to B Metro Mirror volume pairs. A full volume copy is bypassed and only changes are copied from the A volumes to the B volumes.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P -type mmir -mode nocp -incrementalresync override 2100-2107:2100-2107

The following example represents the output:

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.

See "Creating a Metro Mirror relationship" on page 569 for more information.

20. At local site, start the incremental resynchronization with the initialization option on the A volumes in Metro Mirror relationships. The first pass of copying data between the A to B volumes starts (without a full copy). The B to C volumes data copying can also be in the first pass resulting from the failback operation.

Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM.2107-1301261 -type mmir -mode nocp -incrementalresync enable 2100-2107:2100-2107

The following example represents the output:

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.

See "Creating a Metro Mirror relationship" on page 569 for more information.

- 21. Wait until the first pass of the A to B volume pairs to reach full duplex: You can monitor this activity by querying the status of the A to B volumes. As soon as the number of out-of-sync tracks reaches zero, all data has been copied and the data on the A to B volumes is equal. Global Mirror processing starts to form consistency groups when the status of the A to B volumes is full duplex. See "Viewing information about Metro Mirror relationships" on page 578 for more information.
- 22. Start the Global Mirror session at the intermediate site: Enter the mkgmir command at the dscli command prompt with the following parameters and variables (from the local site):

mkgmir -dev IBM.2107-75ALA2P -lss 07 -session 31

The following example represents the output:

CMUC00162I mkgmir: Global Mirror for session 31 successfully started.

See "Starting Global Mirror processing" on page 590 for more information.

Your original configuration is restored.

Chapter 10. Command-line interface scenarios

These scenarios describe some typical configuration and configuration management tasks. You can use them as models for writing your own scripts.

Modifying fixed block volume groups

This scenario describes how to modify fixed block storage within a storage image.

To modify fixed block volume groups, you must have the command-line interface prompt, and you must be connected to a storage image that will be used for open systems host system storage.

Adding volumes to a volume group and removing volumes from a volume group are typical storage management tasks. The volumes that are added to a volume group can be "unassigned" to a volume group, or they can be volumes that are assigned to a volume group but you want to move them to a different volume group. In either case, you are responsible for managing how the volumes are allocated to volume groups and how the volumes are reserved for future allocation. It is better that you maintain "unassigned" volumes in a volume group that is not accessible by any host system, controlling the accessibility of volumes that are reserved for future allocation.

You can assign a fixed block volume to multiple volume groups. This might be necessary for some host system applications. However, damage to volume data can occur if a volume is accessed by different host systems using different file management systems. To assign a fixed block volume to multiple volume groups, complete the following steps:

1. Find the fixed block volumes that are to be assigned to a volume group using the following command.

dscli> lsfbvol -dev ID -datatype 512 | 520p | 520u -extpool ID

The command creates a list of all volumes of the specified volume type within the specified extent pool. It includes only the volumes that are contained by the specified storage image.

 Retrieve the current volume group volume map using the following command. dscli> showvolgrp -dev ID volume_group_ID

The command creates a list of volumes that are assigned to the target volume group.

3. Modify the volume group using the following command.

dscli> chvolgrp -dev *ID* -action add | remove | replace -volume *ID*, *ID*, ..., *ID* volume_group_*ID* To add or remove volumes, you can add or remove volume IDs to the list. This command applies the updated volume ID list.

Deleting data storage configurations

This section describes how you can delete or remove fixed block or count key data storage within a storage image by using the command-line interface. This applies to entire configurations and not just the removal of a volume or volume group.

Before you begin, you must be logged into the DS CLI in interactive command mode. You must also be connected to a storage image that is used for open systems host system storage.

Deleting data storage configurations involves the following steps:

- 1. Remove host access to the volumes that will be removed.
 - For fixed block storage , the SCSI host port IDs must be removed.

- For count key data storage, the CKD volumes are automatically removed from the FICON/ESCON-All volume group ID (10) when the CKD volumes are deleted.
- 2. Remove the volume groups.
 - For fixed block storage, if all of the fixed block volumes are being removed, the associated volume groups must be removed.
 - The FICON/ESCON-All volume group is automatically removed when the last CKD volume is removed.
- 3. Remove the volumes. This applies to both fixed block and count key data storage.
- 4. Remove the logical control units (CKD only).

Note: Logical subsystems (LSS) are automatically removed when the fixed block volumes are removed.

- 5. Remove the ranks.
- 6. Remove the arrays.
- 7. Remove the extent pools.

When all the steps have been completed, the array sites that have been freed are designated as unassigned. They can be redefined to make new fixed block or CKD storage resources.

Deleting a fixed block data storage configuration

Complete this task to delete a fixed block data storage configuration.

To delete fixed block data storage, you must have the command-line interface prompt, and you must be connected to a storage image that contains configured storage.

Deleting a storage configuration involves several steps that systematically remove host access to the data storage, and then removes the storage elements (arrays, ranks, extent pools, volumes, and volume groups) to restore the physical resource to an "equivalent to new" state.

To delete fixed block data storage, complete the following steps:

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

- 1. Remove host access to the volumes that are to be removed. This task requires the issuance of the **lshostconnect** command and the **rmhostconnect** command.
 - **a.** Issue the **lshostconnect** command to display a list of SCSI host port IDs that are associated with the storage to be removed. Enter the **lshostconnect** command at the dscli command prompt as follows:

dscli> lshostconnect -dev IBM.2107-75FA120 -l -portgrp 1

Notes:

- 1) The **-portgrp** *port_grp_number* parameter is used to only list those port IDs that are associated with a port group number that you assigned when you originally created the host connection.
- 2) The -l parameter is used to generate the detailed status report for each host connection.
- b. Issue the **rmhostconnect** command to delete the SCSI host port IDs that are associated with the storage volumes to be removed. Enter the **rmhostconnect** at the dscli command prompt as follows: dscli> rmhostconnect -dev IBM.2107-75FA120 1

Notes:

1) The host_connect_ID parameter (1 in the command example) is required and is a unique identifier (0 - 65 534) within the scope of a storage image.

- 2) A message is displayed with a request that you confirm the deletion of the host connection.
- **2**. Find the volume groups and volume group storage maps by issuing the **lsvolgrp** and **showvolgrp** commands.
 - a. Issue the **lsvolgrp** command to display a list of defined volume group IDs and their characteristics. Enter the **lsvolgrp** command at the dscli command prompt as follows: dscli> lsvolgrp -dev IBM.2107-75FA120 -1
 - b. Issue the showvolgrp command to display the detailed properties of the volume group that you want to delete. Enter the showvolgrp command at the dscli command prompt as follows: dscli> showvolgrp -dev IBM.2107-75FA120 -lunmap V1001

Repeat the **showvolgrp** command for each volume group you want to delete.

Note: The *Volume_Group_ID* (V1001) parameter is required. The shortened form is allowed when you designate the **-dev** parameter.

- **c**. Copy the list of volumes within the volume group for later use when you analyze which volumes that you want to remove.
- Remove the volume groups, as a means to remove volume access by host systems, by issuing the rmvolgrp command. Enter the rmvolgrp command at the dscli command prompt as follows: dscli> rmvolgrp -dev IBM.2107-75FA120 V123-V125

Notes:

- a. All volume groups that are specified for deletion must belong to the same storage unit.
- b. The *Volume_Group_ID* parameter (V123-V125 in the example) is required. If you designate the **-dev** parameter, the shortened version of the ID is allowed.
- **c.** The example command shows a range of volume group IDs. If you have another volume group or another volume group range, you must add a blank between the designations (for example, V123-V125 V130-V133 V135).
- d. A message is displayed for each deleted volume group ID or range of volume group IDs. The message requests that you confirm the deletion.
- 4. Remove the fixed block volumes by issuing the **rmfbvol** command. This action enables the removal of the associated ranks, arrays, and extent pools. Enter the **rmfbvol** command at the dscli command prompt as follows:

dscli> rmfbvol -dev IBM.2107-75FA120 0100 0101

Notes:

- **a**. The associated logical subsystem (LSS) is automatically removed when the last volume that is contained by the LSS is removed.
- b. The *Volume_ID* parameter (represented by 0100 0101 in the example) is required when you issue the **rmfbvol** command. If you designate the **-dev** parameter, the shortened version of the ID is allowed.
- **c**. A message is displayed for each volume that is deleted. The message requests that you confirm the deletion.
- 5. Remove the ranks by issuing the lsrank and rmrank commands.
 - **a.** Issue the **lsrank** command to display a list of rank IDs to be removed. Use the command parameters to develop a selective list of rank IDs. Enter the **lsrank** command at the dscli prompt as follows:

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

Note: Rank IDs that indicate extents used = 0 are eligible to be removed. If extents used are greater than 0 then rank segments are currently assigned to existing volume IDs.

b. Issue the **rmrank** command to remove the ranks that are assigned to the arrays. Enter the **rmrank** command at the dscli prompt as follows:

dscli> rmrank -dev IBM.2107-75FA120 R23

Notes:

- 1) The *rank_ID* parameter (R23 in the example) is required. If you designate the **-dev** parameter, the shortened version of the ID is allowed.
- 2) You must remove the ranks before you can remove the arrays and extent pools.
- **3**) The processing time that is associated with the **rmrank** command can be lengthy and might inhibit your use of the array on which this command is being processed.
- 4) When the **rmrank** command is issued, the following processing occurs:
 - The rank is unassigned from the array.
 - The rank is removed. When this process 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 processing, the array cannot be removed or assigned to another rank. Also, until this process is fully completed, the rank is listed as assigned to the array from which it is has been removed.
 - 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. When you no longer see the rank that is assigned to the array from which you removed it, the remove rank process is complete.
- 6. Remove the arrays by issuing the lsarray and rmarray commands.
 - **a**. Issue the **lsarray** command to obtain a list of array IDs to be removed. Enter the **lsarray** command at the dscli prompt as follows:

dscli> lsarray -dev IBM.2107-75FA120 -state unassigned

Notes:

- 1) The **-state unassigned** parameter allows you to narrow your list to just the array IDs that are not assigned to a rank ID.
- 2) If you issue the lsarray command without using the -state parameter, a list of arrays that have a state of unavailable can appear. This state is a good indication that the ranks have not been removed and that the drives are still formatting. You must wait until the ranks have been removed and the drives have been formatted before you can proceed.
- **3)** Proceed to the next step (remove arrays) only after all the associated arrays are displayed with a state of unassigned.
- b. Issue the **rmarray** command to delete the unassigned arrays so that the array sites can be redefined as new arrays. Enter the **rmarray** command at the dscli command prompt as follows: dscli> rmarray -dev IBM.2107-75FA120 A44-A48 A51

Notes:

- 1) The example command displays the use of a range of array IDs plus one additional array ID.(A44-A48 A51). A range of arrays requires the use of a hyphen and a space between the next array or another range of arrays.
- 2) A message is displayed for each array being deleted that requests your confirmation before processing.
- 7. Remove the extent pools by issuing the **lsextpool** and **rmextpool** commands.
 - **a**. Issue the **lsextpool** command to obtain a list of extent pool IDs to be removed. Enter the **lsextpool** command at the dscli command prompt as follows:

dscli> lsextpool -dev IBM.2107-75FA120 -l -stgtype fb

Note:

- The **-stgtype fb** parameter allows you to narrow the list so that it displays only those extent pools that are assigned for use with fixed block volumes.
- Extent pool IDs that indicate assigned ranks = 0 are eligible to be removed. If the assigned ranks are greater than 0, the extent pool potentially contains assigned storage volumes. The rank indicator must be 0 before you can remove the extent pool.
- b. Issue the rmextpool command to delete extent pool IDs that do not contain assigned rank IDs. Enter the rmextpool command at the dscli command prompt as follows: dscli> rmextpool -dev IBM.2107-75FA120 P21-P25 P30

Notes:

- 1) All rank assignments must be deleted before the extent pool can be deleted.
- **2)** The example command displays the use of a range of extent pool IDs plus one additional extent pool ID (P21-P25 P30). A range of extent pool IDs requires the use of a hyphen and a space between the next extent pool ID or next range of extent pool IDs.

Deleting a count key data storage configuration

Complete this task to delete a count key data storage (CKD) configuration.

To delete CKD storage, you must have the command-line interface prompt, and you must be connected to a storage image that contains configured storage.

Deleting a CKD storage configuration starts with the removal of the CKD volumes and proceeds with the removal of each of the other elements (ranks, arrays, and extent pools) to restore the physical resource to an "equivalent to new" state.

Note: There is no reason to remove the volume groups because the internal code automatically assigns and unassigns CKD volumes to the FICON/ESCON-All volume group ID (10).

To delete CKD storage, complete the following steps:

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Remove the CKD volumes by issuing the **lsckdvol** and **rmckdvol** commands.

a. Issue the **lsckdvol** command to display a list of CKD volume IDs. Analyze the list to determine which IDs can be removed. Enter the **lsckdvol** command at the dscli command prompt as follows: dscli>lsckdvol -dev IBM.2107-75FA120 -lcu 00 -l

Notes:

- 1) You can narrow the list of volume IDs for the designated storage image by using the supported parameters of the **lsckdvol** command.
- 2) The example displays the use of the **-lcu** parameter with a value of 00. Logical control unit (LCU) values are in the range 00 FE for the DS8000 and in the range 00 1E for the DS6000. You must specify a specific LCU; otherwise, the entire storage unit is queried, which results in a longer processing time.
- b. Issue the **rmckdvol** command to delete volumes. This action enables the removal of the associated ranks, arrays, and extent pools. Enter the **rmckdvol** command at the dscli command prompt as follows:

dscli>rmckdvol -dev IBM.2107-75FA120 0100 0101

Note:

• The *Volume_ID* parameter (represented by the values 0100 0101 in the command example) is required when you issue the **rmckdvol** command.

- A message is displayed for each volume that is being deleted. The message requests that you confirm the deletion.
- Issue the rmlcu command to delete LCUs so that the address groups can be redefined for use with fixed block or CKD volumes. Enter the rmlcu command at the dscli command prompt as follows: dscli>rmlcu -dev IBM.2107-75FA120 00-03 08

Note: The example command displays the use of a range of LCU IDs plus one additional LCU ID (00-03 08). A range of LCU IDs requires the use of a hyphen. If you add an additional LCU ID or a range of LCU IDs, you must allow a space between the next LCU ID or another range of LCU IDs.

- 3. Remove the ranks by issuing the **lsrank** and **rmrank** commands.
 - **a**. Issue the **lsrank** command to display a list of rank IDs to be removed. Use the **lsrank** command parameters to develop a selective list of rank IDs. Enter the **lsrank** command at the dscli command prompt as follows:

```
dscli>lsrank -dev IBM.2107-75FA120 -1
```

Note: Rank IDs that indicate extents used = 0 are eligible to be removed. If the displayed value for extents used is greater than 0, it indicates that the ranks are currently assigned to existing volume IDs.

b. Issue the **rmrank** command to remove the ranks that are assigned to the arrays. Enter the **rmrank** command at the dscli prompt as follows:

dscli>rmrank -dev IBM.2107-75FA120 R23

Notes:

- 1) You must remove the ranks before you can remove the arrays and extent pools.
- 2) The processing time that is associated with the **rmrank** command can be lengthy and might inhibit your use of the array on which this command is being processed.
- 3) 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 processing the array cannot be removed or assigned to another rank. Also, until this process is fully completed, the rank is listed as assigned to the array from which it is has been removed.
 - You can check the progress of the **rmrank** command by logging onto another session of DS CLI. Issue the **lsarray** command against the storage image where the rank or ranks are being deleted. When you no longer see the rank that is assigned to the array from which you removed it, the remove rank process is complete.
- 4. Remove the arrays by issuing the **lsarray** and **rmarray** commands.
 - **a**. Issue the **lsarray** command to obtain a list of array IDs to be removed. Enter the **lsarray** command at the dscli prompt as follows:

dscli>lsarray -dev IBM.2107-75FA120 -state unassigned

Notes:

- 1) The **-state unassigned** parameter allows you to narrow your list to just the array IDs that are not assigned to a rank ID.
- 2) If you issue the **lsarray** command without using the **-state** parameter, it is possible you will see a list of arrays that have a state of unavailable. This is a good indication that the ranks have not been removed and that the drives are still formatting. You must wait until the ranks have been removed and the drives have been formatted before you can proceed.

Proceed to the next step (remove arrays) only after all the associated arrays are displayed with a state of unassigned.

b. Issue the **rmarray** command to delete the unassigned arrays so that the array sites can be redefined as new arrays. Enter the **rmarray** command at the dscli command prompt as follows: dscli>rmarray -dev IBM.2107-75FA120 A44-A48 A51

Notes:

- 1) The example command displays the use of a range of array IDs plus one additional array ID (A44-A48 A51). A range of arrays requires the use of a hyphen and a space between the next array or another range of arrays.
- 2) A message is displayed for each array being deleted that requests your confirmation before processing.
- 5. Remove the extent pools by issuing the **lsextpool** and **rmextpool** commands.
 - **a**. Issue the **lsextpool** command to obtain a list of extent pool IDs to be removed. Enter the **lsextpool** command at the dscli command prompt as follows:

dscli>lsextpool -dev IBM.2107-75FA120 -stgtype fb -l

Notes:

- 1) Use the **-stgtype fb** parameter to narrow the list so that it displays only those extent pools that are assigned for use with fixed block volumes.
- 2) Extent pool IDs that indicate assigned ranks = 0 are eligible to be removed. If the value for assigned ranks is greater than 0, the extent pool potentially contains assigned storage volumes. The rank indicator must be 0 before you can remove the extent pool.
- b. Issue the **rmextpool** command to delete extent pool IDs that do not contain assigned rank IDs. Enter the **rmextpool** command at the dscli command prompt as follows:

dscli>rmextpool -dev IBM.2107-75FA120 P21-P25 P30

Notes:

- 1) All rank assignments must be deleted before the extent pool can be deleted.
- 2) The example command displays the use of a range of extent pool IDs plus one additional extent pool ID (P21-P25 P30). A range of extent pool IDs requires the use of a hyphen. When you add an additional extent pool ID or another range of extent pool IDs, you must put a space between the current extent pool ID value and the next extent pool ID value.

Processing remote FlashCopy (inband) transactions

This scenario describes how to successfully process remote FlashCopy (formerly known as inband FlashCopy) transactions. These transactions can be processed by using the DS CLI remote FlashCopy commands. These transactions cannot be managed through the GUI.

You must be logged in to the DS CLI in interactive command mode.

Remote FlashCopy commands are issued to a source volume of a remote mirror and copy volume pair on a local storage unit. This process enables a FlashCopy pair to be established at the remote site and eliminates the need for a network connection to the remote site solely for the management of FlashCopy. The following steps are based on an example that uses the sites LSS 22 and LSS 2A.

Note: You can enter the commands that are described in the steps either for a DS8000 model or for a DS6000 model. The storage image ID for the DS6000 model is different.

1. You must determine which volumes are available for use and then establish a remote mirror and copy path between LSS 22 and LSS 2A.

a. Enter the **lsavailpprcport** command to obtain a report that lists which volumes are available for use.

dscli> lsavailpprcport -dev IBM.2107-1300861 -remotedev IBM.2107-1300871 -remotewwnn 5005076303FFC03D 22:2A

The following report is generated.

Local port	Attached port	Туре
10030	I0031	FCP
10031	I0030	FCP
I0100	I0101	FCP
I0101	I0100	FCP

b. Enter the **mkpprcpath** command to establish the remote mirror and copy path between LSS 22 and LSS 2A.

dscli> mkpprcpath -dev IBM.2107-1300861 -remotedev IBM.2107-1300871 -remotewwnn 5005076303FFC03D

-srclss 22 -tgtlss 2A I0030:I0031 I0100:I0101

The following confirmation is displayed if your command input is correct.

 ${\tt CMUC00149I}$ mkpprcpath: Remote Mirror and Copy path 22:2A successfully established.

- 2. Enter the mkpprc command to establish a remote mirror and copy pair (2200 to 2A00).
 - dscli> mkpprc -dev IBM.2107-1300861 -remotedev IBM.2107-1300871 -type mmir 2200:2A00

The following confirmation is displayed if your command input is correct.

CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship 2200:2A00 successfully created.

3. Enter the **mkremoteflash** command to use LSS 22 on the local site as a conduit LSS for new remote FlashCopy relationships on the remote storage unit. These new relationships use volume 2A00 as their source. The target can be any other volume on the remote storage unit (in this scenario, 2A01).

dscli> mkremoteflash -dev IBM.2107-1300871 -conduit IBM.2107-1300861/22 -record 2A00:2A01

The following confirmation is displayed if your command input is correct.

CMUC00173I mkremoteflash: Remote FlashCopy volume pair 2A00:2A01 successfully created. Use the lsremoteflash command to determine copy completion.

4. Enter the **resyncremoteflash** command because the remote FlashCopy relationship (2A00:2A01) was created with the **-record** parameter.

dscli> resyncremoteflash -dev IBM.2107-1300871 -conduit IBM.2107-1300861/22 -record 2A00:2A01

The following confirmation is displayed if your command input is correct.

CMUC00175I resyncremoteflash: Remote FlashCopy volume pair 2A00:2A01 successfully resynched. Use the lsremoteflash command to determine copy completion.

5. Enter the lsremoteflash command to verify that the transaction processed as you intended.

dscli> lsremoteflash -dev IBM.2107-1300871 -conduit IBM.2107-1300861/22 2A00:2A01

The following report is displayed if your command input is correct.

ID	SrcLSS	Sequence Num	ActiveCopy	Recording
2A00:2A01	2A	0	Disabled	Enabled

Persistent	Revertible	SourceWrite Enabled	TargetWrite Enabled	Background Copy
Enabled	Disabled	Disabled	Disabled	Enabled

Metro Mirror test scenario: failback operation from local to remote site

This scenario describes the steps required to test the failover and failback procedures in which a failback is done from the local site to the remote site.

This test allows you to start a test application on the remote volumes. Then, after the test is complete, resynchronize the remote volumes from the local (production) volumes by copying only changed tracks.

Assume the following cases for this scenario:

- Production is running at Site A (the local site).
- You have simulated a disaster by disabling the links between the local and remote storage units.

Complete these steps for the failover and failback test scenario. (The parameters and values included in this scenario are examples.)

Note: You can issue the commands that are described in the steps either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

1. Freeze updates to the primary (A) volumes in Metro Mirror relationships across the affected LSSs. This process ensures that the secondary (B) volumes will be consistent at the time of the freeze. (One command per LSS is required.) Enter the **freezepprc** command at the dscli command prompt with the following parameters and variables:

freezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07:12

The following example represents the output:

CMUC00161W freezepprc: Remote Mirror and Copy consistency group 07:12 successfully created.

Following the freeze action, the following results occur:

- I/O processing to the Metro Mirror volume pairs is temporarily queued during the time that updates are frozen.
- The volume pairs that are associated with the source and target LSSs are suspended. During this time, updates are collected using the change recording feature on the Site A volumes.
- The established paths between the LSS pairs are disabled.

2. Resume operations following a freeze.

Issue the **unfreezepprc** command to allow I/O activity to resume for the specified volume pairs. Enter the **unfreezepprc** command at the dscli command prompt with the following parameters and variables:

Note: This action is sometimes referred to as a *thaw* operation. dscli> unfreezepprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P 07:12 The following example represents the output: CMUC00198I unfreezepprc: Remote Mirror and Copy pair 07:12 successfully thawed.

3. **At Site B (remote site), issue a failover command to the B to A volume pairs.** Enter the **failoverpprc** command at the dscli command prompt with the following parameters and variables:

dscli> failoverpprc -dev IBM.2107-75ALA2P -remotedev IBM.2107-130165X -type mmir 1200-125f:1a00-1a5f

The following example represents the output:

CMUC00196I failoverpprc: Remote Mirror and Copy pair 1200:1A00 successfully reversed. CMUC00196I failoverpprc: Remote Mirror and Copy pair 1201:1A01 successfully reversed.

When this command processes, the following results occur:

- The B volumes become suspended primary volumes. Updates are collected using the change recording feature on the volumes.
- The A volumes are suspended primary volumes.
- 4. Allow test I/O to start at Site B.
- 5. When testing is complete, complete the following steps:
 - a. Quiesce test I/O at Site B (remote site).
 - b. Enable the remote mirror and copy links between the storage units across the two sites. (The paths will not reestablish automatically.)
 - c. Reestablish paths between the local and remote site LSSs that contain the Metro Mirror volume pairs. Enter the mkpprcpath command at the dscli command prompt with the following parameters and variables:

```
dscli> mkpprcpath -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P
-remotewwnn 5005076303FFC550 -srclss 07 -tgtlss 12 -consistgrp
I0102:I0031 I0002:I0102
```

The following example represents the output:

CMUC00149I mkpprcpath: Remote Mirror and Copy path 07:12 successfully established.

6. At the local site, issue a failback command to the A to B volume pairs. Enter the failbackpprc command at the dscli command prompt with the following parameters and variables

```
dscli> failbackpprc -dev IBM.2107-130165X -remotedev IBM.2107-75ALA2P
-type mmir 1a00-1a5f:1200-125f
```

The following example represents the output:

CMUC00197I failbackpprc: Remote Mirror and Copy pair 1200:1A00 successfully failed back. CMUC00197I failbackpprc: Remote Mirror and Copy pair 1201:1A01 successfully failed back.

When this command processes, the following results occur:

- Updates that are made to the volumes at Site B are recorded with the change recording feature. Changed tracks of data are copied from the Site A volumes to Site B volumes.
- When the copy process is complete, the Site A volumes will be synchronized with the Site B volumes.
- 7. Production I/O continues to the A volumes.

Allowed remote mirror and copy volume pair conversions

This topic describes allowed volume pair conversions using the remote mirror and copy function.

You can convert remote mirror and copy volume pairs between copy modes. For example, you can convert volume pairs in Global Copy to Metro Mirror mode and vice versa. If you create a Global Copy volume pair where the source volume was associated with a DS8000 or a DS6000 machine type and the target volume was associated with an ESS 2105 Model 800 or 750, you can convert that volume pair to Metro Mirror mode.

Before you establish remote mirror and copy volume pairs, logical paths must be established between the source and target logical subsystem (LSS). I/O ports must be available and configured before you can establish paths between the source and target LSSs. Each LSS with source volumes requires at least one path to be established to the LSS that holds the target volumes.

Note: You can issue the commands that are described below either for a DS8000 model or for a DS6000 model, but for the DS6000 model the storage image ID is different.

The following **mkpprcpath** command establishes remote mirror and copy paths:

```
dscli> mkpprcpath -dev storage_image_ID -remotedev storage_image_ID
-remotewwnn wwnn -srclss source_LSS_ID -tgtlss target_LSS_ID
source_port_ID:target_port_ID
```

Example

dscli> mkpprcpath -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -srclss 01 -tgtlss 00 -remotewwnn 12341234000A000F I1A10:I2A20

Convert Metro Mirror volume pairs to Global Copy mode

You can convert a Metro Mirror volume pair to Global Copy mode. For example, because Global Copy can operate at very long distances, well beyond the 300 km (maximum supported distance for Metro Mirror), you might want to convert some Metro Mirror volume pairs, which contain less critical application data, to Global Copy mode.

The following **mkpprc** command converts Metro Mirror volume pairs Global Copy mode:

dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID
-type gcp SourceVolumeID:TargetVolumeID

Example

```
dscli> mkpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 0100:0100 -type gcp 0101:0101 0102:0102 0103:0103
```

Convert Global Copy volume pairs to Metro Mirror mode

There are two common situations when you would convert a Global Copy volume pair to a Metro Mirror mode:

- You have used Global Copy to complete the bulk transfer of data in the creation of many copy pairs, and you now want to convert some or all of those pairs to Metro Mirror mode. This process resynchronizes the volume pairs by copying all changed data from the source volumes to the target volumes.
- You have Global Copy pairs for which you want to make FlashCopy backups on the remote site. You convert the pairs temporarily to synchronous mode to obtain a point-in-time consistent copy.

The following **mkpprc** command converts Global Copy volume pairs to Metro Mirror mode:

dscli> mkpprc -dev storage_image_ID -remotedev storage_image_ID
-type mmir SourceVolumeID:TargetVolumeID

Example

```
dscli> mkpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 0100:0100 -type mmir 0101:0101 0102:0102 0103:0103
```

Resource Groups

Restrictions for using resource groups for Copy Services scope limiting

A resource group is a policy object containing policy attributes that apply to the resources that are associated with the resource group. The object classes that are associated with a resource group are count key data (CKD) logical volumes, fixed block logical volumes, CKD LCUs and fixed block LSSs. Each object in these object classes has a resource group attribute that specifies the single resource group that the object is assigned to.

Each resource group has a resource group label (RGL) that provides a unique name for the resource group. Certain policy attributes in the resource group specify a resource scope (RS), and each user session has an assigned user resource scope (URS). A RS or URS is a pattern that can be matched to the RGLs in all of the resource groups. A resource group (RS) can select a set of RGs whose associated resources are within the scope of the RS.

The URS allows access only to the resources that the user ID has permission to issue Copy Services requests to. 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 the LSS or LCU. A URS can be set to either a wild card (*), or a text string with no wildcard.

When there is a policy attribute that specifies a RS, the RS allows access only to the resources that the resources within this resource group are allowed to relate to according to the specified policy. The attributes of a policy might involve resources that are on other devices (storage images) when Copy Services relationships are involved with the policy. A RS can be set to either a wildcard (*), a text string with no wildcard, or the null string.

A URS or RS set to wildcard (*) matches any RGL. A URS or RS set to a string value matches an RGL that has the same string value. A RS set to null does not match any RGL. Additionally, a URS or an RS other than a source resource scope (SRS, which is the scope of a primary/source logical volumes that a secondary/target logical volume might have) that is set to a string value also matches the RGL **PUBLIC**.

Notes:

- 1. When resource groups are used to partition the resources of the storage facility between tenants of a multi-tenancy environment, this usually means that all of the resources assigned to a given tenant are associated with a single resource group that is configured for the tenant's resources. Such a resource group for tenant *Tenant1* might typically have an RGL=*Tenant1*, CSGRS=*Tenant1*, PGRS=*Tenant1*, GMMastersAllowed=Tenant1_Sessions, and GMSessionsAllowed=Tenant1_Sessions where Tenant1_Sessions would specify one or more Global Mirror session numbers that are to be used exclusively by the tenant. When a tenant is enabled to use Copy Services, any LCUs or LSSs associated with the tenant's logical volumes must also be assigned to the tenant's resource group along with the tenant's logical volumes to have the policies of the resource group have the desired effect.
- 2. Resource Group 0 is predefined and cannot be created, deleted, or modified. Resource group 0 is the default resource group used when resources are created. The policies in resource group 0 are defined to emulate the behavior of a storage subsystem that does not have support for resource groups.

Some restrictions exist when using resource groups and user resource scopes to manage volumes. One restriction is accomplished by having the system administrator set the RG label (RGL) of the volumes/LSSs whose copy services are to be managed, to that of the URS assigned to the users who are managing them. The system administrator can still manage these volumes, if needed.

Another restriction involves setting the remaining RG settings, particularly the CS global scope that affects all copy services functions between volumes associated with the RG. These limitations are active no matter who sends the command.

There are also existing restrictions for how hosts can access copy services on a unit, including HCDS for CKD hosts. For example, the only RG setting that modifies how CKD hosts can interact with a unit is using copy services pass-through limiting. A CKD volume can be designated as a pass-through volume for data being moved to a fixed block (FB) volume using copy services.

Notes:

- 1. Only exclusive multi-managed groups are supported. For example, there are no sub-managed groups allowed within each managed group.
- 2. Only copy services is restricted between managed groups. Logical configuration can cross between managed group boundaries. For this reason, managed group user IDs should be assigned only copy services or monitor operator authorization. Caution is recommended, since the managed group user can still copy a company volume to another company's volume on the same unit accidentally.
- **3**. The system administrator sets, for a specific managed group, each users URS, the RGL, the CS global scope, and the CS pass-through scope to the same value. These actions ensure the RG completely protects the volumes being managed. For example, your human resources department might need to keep things separate for security or HIPPAA compliance requirements.

Using Resource Groups

This scenario describes how you can use Resource Groups to manage volumes on multiple machines.

Situation

You are the network administrator for a company and are in charge of all of its machines. Within your company, there are several subgroups that need their own space on your machines.

Group_1 is running a seating assignment database for World Cup hosting sites. They need a lot of space for a short amount of time. They use the DS CLI to manage when the FlashCopy relationships are formed. They request the following hardware:

- Machine 1 (located on site)
 - Volumes 0000-000f will be used as development volumes on which hosts perform I/O operations.
 - Volumes 0010-001f will be used for nightly FlashCopys of the development volumes from a DS CLI running on a host machine that Group_1 controls. Group_1 can control the backup if they choose to do so.
- Machine 2 (located on site)
 - Volumes 0000-000f are synchronous PPRC (Metro Mirror) targets of Machine 1's volumes of the same ID.
 - Volumes 0010-001f are in band FlashCopy target volumes that match the source machine.

Group_2 uses the following machines to store client files for a security firm. Due to security audit requirements, data must be held in separate cities in the event of a disaster. They use Tivoli Storage Productivity Center for Replication (TPC-R) to manage these relationships. They request the following hardware:

- Machine 1 (located on site)
 - Volumes 0200-02ff are all serving a giant database which is backed up to a remote site.
- Machine 3 (located 500 miles away)
 - Volumes 0000-00ff are Global Mirror targets of the volumes above.

Group_3 needs a large temporary testing environment to convert a bunch of advertising images from several disparate formats to a common one. They do not need Copy Services overhead, as they want to complete the task as quickly as possible. They request the following hardware:

- Machine 1 (located on site)
 - Volumes 0300-0301 are two large volumes, one of which holds the unprocessed advertising images, the other of which holds the formatted version of the images.

Group_4 asks you to assist in hosting their customer management databases with data mining using z Systems. A FlashCopy will be done on the data that is analyzed at that point in time. When the analysis

is complete and stored, the FlashCopy will be reestablished and the analysis will be redone. This will occur over and over again, so the z Systems will control the FlashCopy with in-band commands. You set up the following volumes to handle this:

- Machine 1 (located on site)
 - Volumes 0400-040f are the base volumes used for the database.
 - Volumes 0410-041f are the FlashCopy targets to be used for data mining.

Prerequisites and assumptions

To prepare for these groups, you set up the following cs_operator user account IDs:

- 1. Group_1 will use this ID in their DS CLI scripts to make their FlashCopy relationships. dscli>mkuser -pw xxxx -group op copy services -scope "g1" g1 cs user
- Group_2, will use this ID in their TPC-R connection. dscli>mkuser -pw xxxx -group op_copy_services -scope "g2" g2_cs_user
- **3**. No user ID will be created for Group 3 or Group 4 since they are not doing Copy Services thorough TPC-R, the DS8000 Storage Management GUI, or the DS CLI.

Configuration details

Next you will create the following resource groups:

- 1. You set the label to *g1* to ensure that only users/hosts with scope *g1* or "*" can issue Copy Services commands to these volumes (all Copy Services commands issued to these volumes will fail if they are not from one of these scopes).
- 2. You set copyglobalscope to *g1* to ensure that the FlashCopy works from/to targets/sources that have a label of *g1*. If this is not set up, it wouldn't as both source and target FlashCopy volumes would not have a matching copyglobalscope/label pairing.
- **3**. You create *RG1* to ensure that the Metro Mirror works between Machine 1 and Machine 2 by setting the target PPRC volumes to *RG1* as well. If this is not setup, the source and target PPRC volumes would not have a matching copyglobalscope/label pairing.
- 4. You set the passglobalscope to *g*1.

Notes:

- **a**. Although it does not affect this particular scenario, it is recommended that you set the passglobalscope equal to the copyglobalscope in all cases.
- b. *g1* was used for copyglobalscope and passglobalscope instead of '*'. '*' would have allowed the relationships to be formed, but it wouldn't have prevented these volumes from affecting volumes or being affected by volumes in other resource groups.
- 5. You issue the following commands:

```
mkresgrp -label "g1" RG1
manageresgrp -ctrl copyglobal -action set -scope "g1" RG1
manageresgrp -ctrl passglobal -action set -scope "g1" RG1
chfbvol -resgrp RG1 0000-000f 0010-001f
```

6. You issue the following commands on Machine 1 and Machine 3:

```
mkresgrp -label "g2" RG2
manageresgrp -ctrl copyglobal -action set -scope "g2" RG2
manageresgrp -ctrl passglobal -action set -scope "g2" RG2
chckdvol -resgrp RG2 0200-02ff
```

7. To prevent any Copy Services commands from being issued to any volumes that are assigned to this scope, you issue the following commands on Machine 1:

```
mkresgrp -label "g3" RG3
manageresgrp -ctrl copyglobal -action set -scope "" RG3
manageresgrp -ctrl passglobal -action set -scope "" RG3
chfbvol -resgrp RG3 0300-0301
```

- 8. Same as RG1, but needs to be issued on only Machine 1.
- **9**. You need to setup HCDS on their z Systems servers to ensure that only the volumes in RG3 are visible to that host.

```
mkresgrp -label "g4" RG4
manageresgrp -ctrl copyglobal -action set -scope "g4" RG3
manageresgrp -ctrl passglobal -action set -scope "g4" RG3
chfbvol -resgrp RG3 0300-0301
```

- **10**. The remaining volumes are not public to these groups, so you set them to RG10 to prevent the groups from issuing Copy Services commands to them. Specifying these volumes as PUBLIC will prevent any named volume from performing Copy Services operations to them, even by administrators.
- 11. You issue the following commands on Machine 1, Machine 2, and Machine 3.

mkresgrp -label "not_public" RG10
manageresgrp -ctrl copyglobal -action set -scope "not_public" RG10
manageresgrp -ctrl passglobal -action set -scope "not_public" RG10
chfbvol -resgrp RG4 < all the rest of the volumes on the box >

Appendix. Archived CLI information

Deprecated CLI information (which includes deprecated commands, parameters, and behavior changes) has been removed from this publication, and can now be found online in the IBM System Storage DS8000 Information Center.

Deprecated command/parameter	Replaced by	Explanation	
setplex		DS6000 only command	
showplex		DS6000 only command	
setdialhome		DS6000 only command	
setsmtp		DS6000 only command	
setsnmp		DS6000 only command	
setsim		DS6000 only command	
setcontactinfo		DS6000 only command	
showcontactinfo		DS6000 only command	
testcallhome		DS6000 only command	
offloadss		DS6000 only command	
mkpe		DS6000 only command	
sendss		DS6000 only command	
sendpe		DS6000 only command	
1sss		DS6000 only command	
1 spe		DS6000 only command	
closeproblem		DS6000 only command	
lsproblem		DS6000 only command	
setrmpw		This command is used with CEC Network Cards, which are no longer supported.	
setnetworkport		This command is used to configure CEC Network Cards, which are no longer supported.	
lsnetworkport		This command is used to configure CEC Network Cards, which are no longer supported.	
shownetworkport		This command is used to configure CEC Network Cards, which are no longer supported.	
setoutput	setenv, showenv	The setoutput command is superseded by the setenv and showenv commands.	
chsi -etautomode on off	<pre>chsi -etautomode all tiered none</pre>	The on value is replaced by the tiered value, and the off value is replaced by the none value.	

Table 18. Deprecated commands table

Notices

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

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive Armonk, NY 10504-1785 U.S.A.

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

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan, Ltd. 3-2-12, Roppongi, Minato-ku, Tokyo 106-8711 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have

been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

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

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

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.

Homologation statement

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

Index

Α

account setup 27 add volumes to session, CLI 584 adding ranks to extent pool, DS CLI 547 address group 523 AIX CLI installation considerations 4 LIBPATH handling 4 AIX LIBPATH 15 alias volumes 282 applykey 137 arrays status 545 arrays, creating (ckd) 534 arrays, creating (FB) 521 arrays, deleting 546 assigning ranks 535 audit report, console 504

В

background copy 611

С

chaccess 117 changing volume states Metro Mirror (DS CLI) 573, 581 chauthpol 63 chckdvol 267 checking command usage 38 chextpool 239 chfbvol 307 chhostconnect 192 chipsec 130 chkeymgr 95 chlcu 257 chlss 301 chpass 64 chpprc 468 chrank 227 chresgrp 382 chsession 495 chsestg 356 space-efficient storage 356 chsi 165 chsp 152 chsu 158 chuser 65 chvolgrp 342 Cisco MDS 9216 Multilayer Fabric Switch 42, 462, 479 ckd creating LCU 536 CKD deleting configurations 685 volume removal 285 CKD configuration error 537 CKD configuration error, correcting 538 ckd volume creating arrays 534

CKD volume rank 535 ckd volumes, creating 537 CLI displaying WWNN 566 pause 572 postinstallation 20 resuming Metro Mirror relationship 571 CLI commands on a command line 35 CLI commands, overview 45 CLI default configuration profile file 21 CLI help 39 CLI history function 38 CLI interactive mode 38 command line presentation 35 commands 119 applykey 137 chaccess 117 chauthpol 63 chckdvol 267 chextpool 239 chfbvol 307 chhostconnect 192 chipsec 130 chkeymgr 95 chlcu 257 chlss 301 chpass 64 chpprc 468 chprc 468 chrank 227 chresgrp 382 chsession 495 chsestg 356 chsi 165 chsp 152 chsu 158 chuser 65 chvolgrp 342 commitflash 392 commitremoteflash 416 cpauthpol 67 diagsi 167 dscli 49 echo 50 failbackpprc 449 failoverpprc 452 flagless parameters 45 freezepprc 469 helpmsg 52 initckdvol 269 initfbvol 309 lsaccess 117 lsaddressgrp 255 lsarray 218 lsarraysite 213 lsauthpol 67 lsavailpprcport 436 lsckdvol 270 lsda 140

commands (continued) lsddm 148 lsextpool 241 lsfbvol 310 lsflash 397 lsframe 142 lsgmir 491 lshba 146 lshostconnect 196 lshosttype 211 lshostvol 201 lsioport 176 lsipsec 131 lsipseccert 132 lskey 138 lskeygrp 96 lskeymgr 99 lslcu 259 lslss 302 lsperfgrp 377 lsperfgrprpt 378 lsperfrescrpt 380 lsportprof 203 lspprc 455 lspprcpath 438 lsrank 228 lsremoteflash 421 lsresgrp 383 lsserver 168 lssession 499 lssestg 359 lssi 170 lsstgencl 144 lssu 159 lsuser 69 lsvolgrp 345 lsvolinit 353 lsvpn 154 manageaccess 119 manageckdvol 277 manageextpool 244 managefbvol 318 managehostconnect 204 managekeygrp 101 managekeymgr 103 managepwfile 70 managereckey 104 manageresgrp 384 mkaliasvol 280 mkarray 222 mkauthpol 72 mkckdvol 282 mkesconpprcpath 442 mkextpool 245 mkfbvol 321 mkflash 402 mkgmir 479 mkhostconnect 205 mkipsec 133 mkipseccert 134 mkkeygrp 105

commands (continued) mkkeymgr 106 mklcu 261 mkpprc 462 mkpprcpath 445 mkrank 233 mkreckey 107 mkremoteflash 426 mkresgrp 388 mksession 497 mksestg 365 mkuser 72 mkvolgrp 346 offloadauditlog 504 overview structure 45 parameters 45 pausegmir 480 pausepprc 470 resumegmir 482 resumepprc 472 resyncflash 393 resyncremoteflash 417 reverseflash 405 revertflash 409 revertremoteflash 430 rmarray 223 rmauthpol 74 rmckdvol 285 rmextpool 246 rmfbvol 324 rmflash 410 rmgmir 483 rmhostconnect 208 rmipsec 135 rmipseccert 136 rmkeygrp 108 rmkeymgr 108 rmlcu 263 rmpprc 474 rmpprcpath 447 rmrank 234 rmreckey 109 rmremoteflash 431 rmresgrp 389 rmsession 503 rmsestg 368 rmuser 75 rmvolgrp 348 setauthpol 75 setenv 53 setflashrevertible 414 setioport 183 setipsec 136 setoutput 56 setremoteflashrevertible 434 setvpn 153 showaccess 122 showarray 224 showarraysite 216 showauthpol 84 showckdvol 286 showenv 58 showextpool 247 showfbvol 327 showgmir 485 showgmircg 493 showgmiroos 494

commands (continued) showhostconnect 209 showioport 184 showkeygrp 110 showkeymgr 114 showlcu 264 showlss 304 showpass 88 showrank 235 showresgrp 389 showsestg 369 showsi 172 showsp 155 showsu 161 showuser 89 showvolgrp 349 testauthpol 90 unfreezeflash 413 unfreezepprc 477 variables 45 ver 59 who 91 whoami 94 comments, sending xi commitflash 392 commitremoteflash 416 concurrent copy modifying timeout value (DS CLI) 574 conduit LSS 691 configuration error 526 configuration profile file 21 configuring Fibre Channel I/O ports 529, 539 configuring I/O port, System i 529, 539 connection ID, caution 192 consistency group 606 inconsistencies 674 modifying timeout value (DS CLI) 575 consistency group parameters, CLI 585 consistency groups creating (DS CLI) 576 console audit report 504 console mode installation, DS CLI 10 controlling CLI page help 39 conventions syntax vii terminology vii typefaces vii conversions Global Copy to Metro Mirror 694 Metro Mirror to Global Copy 694 converting Metro Mirror (DS CLI) 579 volume pairs 579 Copy Services across 2105 and DS6000 581 across 2105 and DS8000 581 correcting path-related configuration error 568 wrong value for remote WWNN 568 wrong value for target LSS ID 568 correcting a CKD configuration error 538

error 526 correcting FlashCopy relationships, commitflash command 609 correcting FlashCopy relationships, revertflash command 608 count key data volumes creating extent pools 532 cpauthpol 67 create FlashCopy relations, Global Mirror 594 create Global Copy pairs 593 create paths, Global Mirror 592 creating ckd volumes 537 consistency group (DS CLI) 570, 576 extent pools, fixed block 519 fixed block volume groups 527 fixed block volumes 523 FlashCopy relationship 555 LUN volumes 524 Metro Mirror volume pair using a 2105 (DS CLI) 582 persistent FlashCopy 556 ranks for CKD volumes 535 SCSI host port connections 530 volume groups for System i 528 creating a rank, DS CLI 522 creating arrays (ckd) 534 creating arrays (FB) 521 creating extent pools, CKD 532 critical heavy mode 462 critical mode setting (DS CLI) 575 critical volume mode 462

correcting fixed block configuration

D

DDM status 544 default CLI settings 21 default directories 35 default report format 56 delete arrays 546 delete extent pools, DS CLI 543 delete LCUs 553 delete ranks 550 deleting FlashCopy relationships 557 paths 568 single volume relationship (DS CLI) 581 volume pair (DS CLI) 573 deleting CKD configurations 689 deleting fixed block configuration 686 delim report format 56 determining available I/O ports 579 diagsi 167 disableautoresync failbackpprc 449 failoverpprc 452 mkpprc 462 resumepprc 472 disaster recovery drain time 614 failback scenario (DS CLI) 577 failover (DS CLI) 578

disaster recovery (continued) failover B volumes to A volumes 605 first pass 614 OOS 614 quiesce your system 614 restoring operations (planned) to the intermediate site 653 restoring operations to the intermediate site 631, 663 planned outage 625 restoring operations to the remote site 635 restoring to the intermediate site 674 disaster recovery, Global Mirror background copy 611 begin the failback recovery process 612 checking status 605 consistency group analysis 606 correcting FlashCopy relationships, commitflash command 609 end Global Mirror processing 604 fast reverse restore 610 Global Copy failback 618 Global Copy failover 617 process overview 603 reestablish FlashCopy relationships, B to C 612 reestablishing paths 615 resume Global Mirror processing 618 resume Global Mirror processing at the A site 618 resynchronize B volumes to A volumes 613 revertflash and FlashCopy relationships 608 disk drive module 544 displaying paths 566 DS CLI console mode installation 10 create paths 567 creating volume pairs 569, 572 default directories 35 determining 579 first time login 35 graphical mode installation 8 I/O ports 579 installation on Linux 10 interactive mode 38 limitations 42 OpenVMS messages 42 script mode 37 session limitation 42 silent mode installation 14 silent mode installation, DS CLI 14 single-shot mode 36 upgrade preparation 16 DS CLI command format 35 DS CLI commands System i 540 DS CLI exit codes 40 DS CLI installation AIX considerations 4 default directories 4 HP Tru64 considerations 4

DS CLI installation (continued) key aspects 4 mounting the installation CD 7 System i considerations 4 DS CLI Tasks Global Mirror add volumes to session 584 create FlashCopy relationships 594 create Global Copy pairs 593 create session 595 end a session 597 end processing, no script 591 end processing, script mode 590 modify tuning parameters 585 pausing processing 589 query Global Mirror session 588 remove FlashCopy relationships 598 remove Global Copy relationships 599 remove paths 600 remove volumes 596 resume processing 589 setup, creating paths 592 start processing 590 view session 588 DS CLI Tasksmodify topology Global Mirror 586 DS CLI uninstall Linux systems 16 dscli command 49 during an unplanned outage 642

Ε

eam (extent allocation method) lsckdvol 270 lsfbvol 310 mkckdvol 282 mkfbvol 321 showckdvol 286 showfbvol 327 EAV (extended addressing volume) mkckdvol 282 echo command 50 encryption 103 encryption key 103 end Global Mirror script mode 590 ending Global Mirror processing 604 exit codes, obtaining 40 extent pool modifications, CLI 542 extent pool names 519 extent pool, adding ranks with DS CLI 547 extent pools, delete with DS CLI 543 extent pools, System i 519

F

failback disaster recovery scenario (DS CLI) 577 local to remote 693 restoring operations (planned) to the intermediate site 653 failback (continued) restoring operations (unplanned) to the intermediate site 663 using -force 653, 663 failback recovery process 612 failbackpprc 449 disableautoresync 449 tgtse 449 failover disaster recovery (DS CLI) 578 during an unplanned outage 642 planned outage restoring operations to the intermediate site 625 restoring operations (planned) to the intermediate site 653 restoring operations (unplanned) to the intermediate site 663 restoring operations to the intermediate site 631, 674 restoring operations to the remote site 635 using -force 653, 663 failoverpprc 452 disableautoresync 452 tgtse 452 fast reverse restore 610 feedback, sending xi Fibre Channel I/O ports configuring 529, 539 fc-al 529, 539 ficon 529, 539 scsi-fcp 529, 539 first pass 614 first time login 35 fixed block deleting configurations 685 fixed block configuration error, correcting 526 fixed block volume creating arrays 521 fixed block volume groups, creating 527 fixed block volume groups, modifying 685 fixed block volumes creating extent pools 519 fixed block volumes, creating 523 FlashCopy background copy 562 creating a FlashCopy relationship 555 deleting relationships 557 inconsistent data on target volumes 674 persistent creating 556 refreshing target volumes 560 reversing relationship 560 revertible option 561 revertible option defined 561 setflashrevertible 561 target volumes on Metro Mirror source 563 removing write inhibits 563 viewing information about FlashCopy relationships 556 FlashCopy relationships, create 594

force parameter 449, 452 freezepprc 469

G

Global Copy creating volume pairs 572 Global Copy failback 618 Global Copy failover 617 Global Copy pairs, create 593 Global Mirror - CLI add volumes to session 584 create FlashCopy relationships 594 create Global Copy pairs 593 create session 595 end processing, no script 591 end processing, script mode 590 end session 597 modify topology 586 modify tuning parameters 585 pausing processing 589 query session 588 remove FlashCopy relationships 598 remove Global Copy pair 599 remove paths 600 remove volumes 596 resume processing 589 setup environment, paths 592 start processing 590 view session 588 Global Mirror processing, ending 604 graphical mode installation, DS CLI 8

Η

help, CLI 39 helpmsg 52 HMC 122 homologation 705 howaccess 122 HP-UX volume restriction 192

I/O ports available 579 i5/OS command format 36 green screen, DS CLI 540 library QDSCLI 540 removing DS CLI directly 19 silent mode CLI installation 14 i5/OS password file directory 70 inband transactions 558, 691 incremental FlashCopy copying updated tracks only 560 incremental resynchronization during a planned outage at an intermediate site 674 incrementalresvnc 462 initckdvol 269 initfbvol 309 installing the CLI 20

interactive mode history function 38 setenv command 38 IPSec connection file format 124 issuing DS CLI commands 540 i5/OS 540

J

Java Virtual Machine Not Found Error 15

L

LCU creation 536 LCU status 552 LCU status, one 552 LCU, modify with CLI 551 LCUs, deleting 553 library ODSCLI 540 limitations remote FlashCopy 558 Linux DS CLI installation, console mode 10 lshostvol command considerations 201 Linux password file directory 70 Linux, uninstall DS CLI 16 lock account 65 logical control unit creation 536 logical control unit properties 552 lsaccess 117 lsaddressgrp 255 lsarray 218 lsarraysite 213 lsauthpol 67 lsavailpprcport 436 lsckdvol eam (extent allocation method) 270 sam (standard allocation method) 270 lsda 140 lsddm 148 lsextpool 241 lsfbvol 310 eam (extent allocation method) 310 tse 310 lsflash 397 tgtse, space-efficient volume 397 lsframe 142 lsgmir 491 lshba 146 lshostconnect 196 lshosttype 211 lshostvol Linux considerations 201 OpenVMS considerations 201 lsioport 176 lsipsec 131 lsipseccert 132 lskey 138 lskeygrp 96 lskeymgr 99 lslcu 259

lslss 302 lsperfgrp 377 lsperfgrprpt 378 lsperfrescrpt 380 lsportprof 203 lspprc 455 lspprcpath 438 command 566 lsrank 228 lsremoteflash 421 tgtse, space-efficient volume 421 lsresgrp 383 lsserver 168 lssession 499 lssestg 359 space-efficient storage 359 lssi 170 command 566 lsstgencl 144 lssu 159 lsuser 69 lsvolgrp 345 lsvolinit 353 lsvpn 154 LUN sizes 321 LUN volumes, creating 524

Μ

manageaccess 119 manageckdvol 277 manageextpool 244 managefbvol 318 managehostconnect 204 managekeygrp 101 managekeymgr 103 managepwfile 70 managereckey 104 manageresgrp 384 message types 42 Metro Mirror consistency groups 570 creating volume pairs 569 FlashCopy target on Metro Mirror source 563 scenario 693 Metro/Global Mirror restoring operations (planned) to the intermediate site 653 restoring operations (unplanned) to the intermediate site 663 setting up 621 mkaliasvol 280 mkarray 222 mkauthpol 72 mkckdvol 282 eam (extent allocation method) 282 EAV (extended addressing volume) 282 tse 282 mkesconpprcpath 442 mkextpool 245 mkfbvol 321 eam (extent allocation method) 321 space-efficient storage 321 mkflash 402

mkflash (continued) tgtse, space-efficient volume 402 mkgmir 479 mkhostconnect 205 mkipsec 133 mkipseccert 134 mkkeygrp 105 mkkeymgr 106 mklcu 261 mkpprc 462 disableautoresync 462 tgtse 462 volume pairs 569, 572 mkpprcpath 445 creating paths 567 mkrank 233 mkreckey 107 mkremoteflash 426 tgtse, space-efficient volume 426 mkresgrp 388 mksession 497 mksestg 365 space-efficient storage 365 mkuser 72 mkvolgrp 346 modify extent pool, CLI 542 modify a rank, CLI 548 modify an LCU, CLI 551 modify topology, CLI 586 modify tuning parameters, CLI 585 modifying concurrent copy timeout value (DS CLI) 574 consistency group timeout (DS CLI) 575 critical mode setting (DS CLI) 575 fixed block volume groups 685 timeout values (DS CLI) 574 z/OS Global Mirror (DS CLI) 575 mounting the DS CLI installation CD 7 multiple LCU status 552 multiple ranks status 549

Ν

new features xiii nonrevertible 608, 609 Novell DS CLI graphical installation 8

0

offloadauditlog 504 OpenVMS DS CLI message considerations 42 Ishostvol command considerations 201 operational limitations 42 OS/400 DS CLI graphical installation 8 os400all 528 os400mask 528 out-of-sync 614

Ρ

password file directories i5/OS 70 LINUX 70 UNIX 70 Windows 70 paths creating 567 deleting 568 displaying 566 lspprcpath 577 mkpprcpath 567 port considerations 567 view existing 577 Paths correcting configuration error 568 pausegmir 480 pausepprc 470 suspending volumes 572 persistent FlashCopy creating 556 planned outage restoring operations to the intermediate site 625 postinstallation CLI 20 postinstallation tasks 20 primary volumes 605 processing tips 3 site recovery 462 cascade 462 incrementalresync 462 multiple volumes 462 volume grouping method 462 publications ordering xi product viii related viii

Q

query Global Mirror session 588 quiesce your system 614

R

RAID 6 creating extent pools 532 fixed block volumes 519 lsarray 218 lsrank 228 LUN volumes 524 showarray 224 rank configuration 522, 535 rank group 1 or 1 535 rank status, one 549 rank, create with DS CLI 522 rank, modify with CLI 548 ranks status 549 ranks, deleting 550 reestablishing remote mirror and copy paths 615 relationship reversing FlashCopy 560 release rank, CLI 548

remote FlashCopy limitations 558 remote FlashCopy transactions 691 Remote Mirror and Copy relationships viewing (DS CLI) 579 remove arrays 546 remove LCUs 553 remove ranks 550 removing uncorrupted incremental backup 563 write inhibits from FlashCopy target volumes 563 removing DS CLI, i5/OS 19 report formats controlling 38 default 56 delim 56 setoutput command 56 stanza 56 XML 56 reserve rank, CLI 548 resource groups using 697 restoring operations (planned) to the intermediate site 653 restoring operations (unplanned) to the intermediate site 663 restoring operations to the intermediate site 625, 631 restoring operations to the remote site 635, 642 resume Global Mirror processing 589 resumegmir 482 resumepprc 472 disableautoresync 472 resuming volumes 571 tgtse 472 resvnc target volumes in FlashCopy relationship 560 resyncflash 393 tgtse, space-efficient volume 393 resynchronize B volumes to A volumes 613 resyncremoteflash 417 tgtse, space-efficient volume 417 reverse a FlashCopy relationship 610 reverseflash 405 tgtse, space-efficient volume 405 revertflash 409 revertremoteflash 430 rmarray 223 rmauthpol 74 rmckdvol 285 rmextpool 246 rmfbvol 324 rmflash 410 tgtse, space-efficient volume 410 rmgmir 483 rmhostconnect 208 rmipsec 135 rmipseccert 136 rmkeygrp 108 rmkeymgr 108 rmlcu 263 rmpprc 474 rmpprcpath 447

rmrank 234 rmreckey 109 rmremoteflash 431 tgtse, space-efficient volume 431 rmresgrp 389 rmsession 503 rmsestg 368 space-efficient storage 368 rmuser 75 rmvolgrp 348

S

script mode 37 script mode, end processing 590 SCSI host port connections creating 530 serial number, obtaining 39 service action report, console 504 service audit log 504 setauthpol 75 setenv 53 setflashrevertible 414 tgtse, space-efficient volume 414 setioport 183 setipsec 136 setoutput 56 setremoteflashrevertible 434 tgtse, space-efficient volume 434 setting up Metro/Global Mirror 621 setup, account 27 setvpn 153 showaccess 122 showarray 224 showarraysite 216 showauthpol 84 showckdvol 286 eam (extent allocation method 286 tse 286 showenv 58 showextpool 247 showfbvol 327 eam (extent allocation method) 327 tse 327 showgmir 485 showgmircg 493 showgmiroos 494 showhostconnect 209 showioport 184 showkeygrp 110 showkeymgr 114 showlcu 264 showlss 304 showpass 88 showrank 235 showresgrp 389 showsestg 369 space-efficient storage 369 showsi 172 showsp 155 showsu 161 showuser 89 showvolgrp 349 single volume relationship deleting (DS CLI) 581

single-shot mode command format, i5/OS 36 command format, Windows 36 space-efficient storage chsestg 356 lsfbvol 310 lssestg 359 mkfbvol 321 mksestg 365 rmsestg 368 showsestg 369 space-efficient volume Isflash 397 lsremoteflash 421 mkflash 402 mkremoteflash 426 resyncflash 393 resyncremoteflash 417 reverseflash 405 rmflash 410 rmremoteflash 431 setflashrevertible 414 setremoteflashrevertible 434 stanza report format 56 starting FlashCopy 562 status information, fixed block volume 310 stop the process Global Mirror 483 storage image ID, obtaining 39 syntax conventions vii System i bypassing uninstaller 19 CLI installation considerations 4 configuring I/O port 529, 539 creating FlashCopy relationships 555 Global Copy pairs 572 Metro Mirror pairs 569 DS CLI commands 540 extent pool considerations 519 host connections 530 interactive mode 38 issuing 540 LUN volumes 524 lunmapping 349 os400all 528 os400mask 528 postinstallation tasks 20 removing CLI, direct method 19 removing CLI, remote method 20 script mode 37 serial number reporting 172 showvolgrp 349 single-shot mode 36 volume group creation external load source consideration 528 multipath consideration 528

T

target volume commit changes 565 discard changes 564

target volume (continued) FlashCopy commit changes 565 discard changes 564 terminology vii testauthpol 90 tgtse failbackpprc 449 failoverpprc 452 lsflash 397 lsremoteflash 421 mkflash 402 mkpprc 462 mkremoteflash 426 resumepprc 472 resyncflash 393 resyncremoteflash 417 reverseflash 405 rmflash 410 rmremoteflash 431 setflashrevertible 414 setremoteflashrevertible 434 thaw operation 477, 625, 631, 635, 642, 653, 663, 693 timeout connection 21 timeout values modifying concurrent copy (DS CLI) 574 consistency group (DS CLI) 574 critical mode enabled (DS CLI) 574 z/OS Global Mirror (DS CLI) 574 topology defined, CLI 586 trademarks 704 tse lsckdvol 270 lsfbvol 310 mkckdvol 282 showckdvol 286 showfbvol 327 space-efficient storage lsckdvol 270 tse, mksestg mkfbvol 321 tuning parameters, CLI 585

U

unassign rank, CLI 548 unattended (silent) mode installation 14 unattended (silent) mode installation, DS CLI 14 unfreezeflash 413 unfreezepprc 477 uninstall.dat 20 uninstall.jar 20 UNIX DS CLI graphical installation 8 UNIX password file directory 70 unlock account 65 user security 119

V

ver 59 view a specific LCU 552 view a specific rank 549 view array sites status 545 view arrays status 545 view LCU status 552 view ranks status 549 view session, CLI 588 viewing existing paths 577 information about FlashCopy relationships 556 Remote Mirror and Copy relationships (DS CLI) 579 volume format processing options 462 volume pair deleting (DS CLI) 573 volume removal, CKD 285 volumes FlashCopy target volumes commit changes 565 discard changes 564 volumes, alias 282

W

websites viii who 91 whoami 94 Windows DS CLI graphical installation 8 Windows password file directory 70 write acceleration restriction 42, 462, 479 write inhibits mkflash command 563 removing from FlashCopy target volumes 563 wwnn displaying for paths 566

X

XML report format 56

Ζ

z/OS Global Mirror modifying critical mode setting (DS CLI) 575 timeout value (DS CLI) 575

Printed in USA

GC27-4212-08

