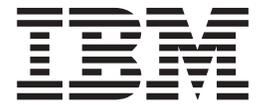


IBM System Storage DS8000



Command-Line Interface User's Guide

Version 1 Release 2

IBM System Storage DS8000



Command-Line Interface User's Guide

Version 1 Release 2

Note:

Before using this information and the product it supports, read the information in the **Safety and environmental notices** and **Notices** sections.

Second Edition (November 2006)

This edition replaces SC26-7916-00.

This edition also applies to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2004, 2006. All rights reserved.**

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

Contents

About this guide vii

Notices and publication information . . . ix

Safety notices	ix
Environmental notices	ix
Product recycling and disposal	ix
Battery return program	x
Conventions used in this guide	xi
Related information	xii
DS8000 series library	xii
Other IBM publications	xiii
Ordering IBM publications	xvii
Web sites	xviii
How to send your comments	xviii

**Summary of changes for SC26-7916-01
IBM System Storage DS8000
Command-Line Interface User's Guide . . . xxi**

**Chapter 1. Introduction to IBM System
Storage DS8000 series 1**

Overview of the DS8000 series models	1
DS8100 (Model 921 and Model 931)	4
DS8300 (Models 922, 932, 9A2, and 9B2)	5
Model comparison	6
IBM System Storage Management Console	7
DS8000 Performance features	8
RAID implementation	9
RAID 5 overview	10
RAID 10 overview	10
DS8000 Interfaces	10
IBM System Storage DS Storage Manager	10
The DS command-line interface	11
DS open application programming interface	11
IBM TotalStorage Productivity Center for Disk	12
IBM TotalStorage Productivity Center for Replication	13
Host systems attachment overview	13
IBM System Storage Resiliency Family	14
Disaster recovery using Copy Services	19
Comparison of licensed functions	21
Logical configuration overview	22

**Chapter 2. Upgrading your system to
use DS CLI 25**

**Chapter 3. Installing and removing the
CLI 27**

Supported operating systems for the DS command-line interface	27
DS CLI operational limitations	27
Preparing for the DS CLI installation	28
Installing the DS CLI using the graphical mode	30

Installing the DS CLI using the console mode	36
Installing the DS CLI using the unattended (silent) mode	42
Correcting the JVM-not-found error	46
Installing DS CLI on an OpenVMS system	47
Preparing to install the DS CLI on an OpenVMS System	48
Installing the DS CLI on your OpenVMS system	48
Using the DS CLI on an OpenVMS system	53
Removing the DS CLI from an OpenVMS system	54
Completing DS CLI postinstallation	54
Using the DS CLI application	55
DS CLI command help	60
Obtaining and interpreting DS CLI exit codes	61
Activating your machine and feature licenses using the DS CLI	63
Setting up user accounts using the DS CLI	64
Creating a default configuration setup with a profile file	65
Removing the DS Command Line Interface	71
Removing the DS CLI from your system using graphical mode	71
Removing the DS CLI using unattended (silent) mode	73
Removing the DS CLI using the console mode	73
Removing the DS CLI from a System i system	74
OpenVMS system integration	76
Enhancing the command console LUN for DS CLI use	76
Enhancing the OpenVMS system messages	77
Enabling OpenVMS to use the DS CLI help	77
Java Runtime Environment considerations for DS CLI	78
Quota considerations for DS CLI	79

**Chapter 4. Configuring and managing
logical storage 81**

Configuring the DS8000 (using DS CLI) for use with the TotalStorage Productivity Center Replication Manager	81
Configuring new fixed block storage using the DS CLI	84
Creating extent pools for fixed block volumes using the DS CLI	85
Creating arrays for fixed block volumes using the DS CLI	86
Creating a rank using the DS CLI	87
Creating fixed block volumes using the DS CLI	88
Creating LUN volumes for System i systems	90
Correcting a fixed block configuration error	92
Creating fixed block volume groups using the DS CLI	92
Creating a volume group for System i systems	94
Configuring fibre-channel I/O ports using the DS CLI	95

Creating SCSI host port connections using DS CLI	96
Configuring new count key data storage using DS CLI	98
Creating count key data extent pools using the DS CLI	99
Creating arrays for CKD volumes using the DS CLI.	101
Creating a rank for CKD volumes using the DS CLI.	102
Creating logical control units for CKD volumes using DS CLI	103
Creating count key data volumes using the DS CLI.	104
Correcting a CKD volume configuration error	106
Configuring fibre-channel I/O ports using the DS CLI	106
Managing your logical storage configuration	107
Using DS CLI commands on i5/OS	108
Modifying an extent pool	110
Viewing extent pool status	111
Viewing extent pool properties and performance metrics	111
Deleting extent pools from a storage configuration	112
Viewing the array disk drive module status	112
Viewing array status	113
Viewing properties for one array	114
Removing arrays from a storage configuration or a rank assignment	114
Adding a rank to an extent pool	116
Modifying a rank	116
Viewing rank status	118
Viewing properties for one rank	118
Correcting a rank-related configuration error	119
Removing ranks from a storage configuration	119
Modifying a logical control unit	120
Viewing logical control unit status	121
Viewing properties for one logical control unit	122
Removing logical control units from a CKD storage configuration.	122

Chapter 5. Copy Services functions 125

FlashCopy functions	125
Creating a FlashCopy relationship	125
Creating a persistent FlashCopy relationship	126
Viewing information about FlashCopy relationships.	126
Deleting FlashCopy relationships.	128
Creating remote FlashCopy transactions	129
Resynchronizing FlashCopy relationships	130
Reversing a FlashCopy relationship	131
Applying the FlashCopy revertible option to existing FlashCopy relationships	132
Starting a background copy of a FlashCopy relationship	133
Preventing write operations on FlashCopy target volumes	134
Creating a FlashCopy target volume on an existing Metro Mirror source volume	135

Discarding changes to FlashCopy target volumes	136
Committing data to FlashCopy target volumes	137
Metro Mirror functions	138
Displaying the status of established paths	138
Displaying the WWNN of a storage unit	138
Creating remote mirror and copy paths.	139
Correcting a path-related configuration error	140
Removing paths	141
Creating a Metro Mirror relationship	142
Creating a Metro Mirror consistency group	143
Resuming a Metro Mirror relationship	144
Pausing a Metro Mirror relationship.	145
Creating a Global Copy relationship.	146
Deleting a Metro Mirror relationship	147
Modifying logical subsystem timeout values	147
Defining a path that has the consistency option enabled	151
Monitoring Remote Mirror and Copy paths	151
Performing a failback recovery operation	152
Performing a failover recovery operation	152
Viewing information about Metro Mirror relationships.	153
Converting Global Copy volume pairs to synchronous.	154
Determining which I/O ports are available for paths	154
Deleting a single volume Metro Mirror relationship	156
Copy Services functions across a 2105 and 2107/242x	157
Creating a Metro Mirror volume pair between a 2107 and a 2105	157
Global Mirror functions	160
Adding volumes to a session (Global Mirror)	160
Modifying the tuning parameters of a Global Mirror session	161
Modifying the topology of a Global Mirror session	162
Viewing a Global Mirror session	164
Querying Global Mirror processing	165
Pausing Global Mirror processing	165
Resuming Global Mirror processing	166
Starting Global Mirror processing	167
Ending Global Mirror processing (script mode)	167
Ending Global Mirror processing (no script)	168
Setting up the Global Mirror Environment.	168
Removing a Global Mirror environment	174

Chapter 6. Recovering from a disaster using Global Mirror. 181

Ending Global Mirror processing when a disaster occurs	182
Checking Global Mirror transaction status in a disaster situation	183
Initiating failover processing for B volumes to A volumes	184
Analyzing and validating the consistency group state	185
Using the revertflash command to correct FlashCopy relationships	187

Using the commitflash command to correct FlashCopy relationships	188
Using fast reverse restore processing to create consistency	189
Waiting for the background copy to complete	190
Reestablishing the FlashCopy relationships between B volumes and C volumes	191
Preparing to reinstate production at the local site	192
Resynchronizing the volumes	193
Querying, quiescing, and querying	194
Reestablishing remote mirror and copy paths (site A to site B)	195
Performing Global Copy failover processing to the A volumes	197
Performing Global Copy failback processing for the A volumes	197
Resuming Global Mirror processing at site A	198

Chapter 7. Recovery scenarios for planned and unplanned outages using Metro/Global Mirror 201

Setting up a Metro/Global Mirror environment	201
Failover and restore operations to the intermediate site during a planned outage	205
Failover and restore operations to the intermediate site during an unplanned outage	212
Failover and restore operations at the remote site during a planned outage	217
Failover and restore operations at the remote site during an unplanned outage	225
Using forced failover and failback during a planned Metro/Global Mirror outage	236
Using forced failover and failback during an unplanned Metro/Global Mirror outage	247
Discarding changes or committing changes to consistency groups	258
Recovery scenario using incremental resynchronization in a Metro/Global Mirror configuration	259

Chapter 8. CLI commands 271

About CLI commands	271
Understanding the syntax diagrams	272
List of commands	273
Command flags	286
Using hexadecimal values	288
Command equivalents	289
Output field descriptions	294
Framework command	305
dscli	305
exit	307
help	307
quit	308
setoutput	309
User account and security commands	312
chpass	312
chuser	313
lsuser	315
managepwfile	316
mkuser	318

rmuser	320
setmpw	321
showpass	321
showuser	322
Application key and version commands	323
applykey	323
lskey	324
ver	325
Physical enclosure information commands	326
lsstgencl	326
lsframe	329
Physical resource information commands	330
lsda	330
lshba	332
lsddm	334
Storage complex configuration commands	338
chsp	338
setvpn	340
showsp	341
Storage unit configuration commands	342
chsu	343
lssu	344
showsu	346
Storage image configuration commands	349
chsi	349
lsserver	350
lssi	352
showsi	354
I/O port and host connection configuration commands	356
Storage image I/O port commands	356
Host connection commands	365
Network ports	388
setnetworkport	388
lsnetworkport	390
shownetworkport	393
Storage configuration commands	395
Array site specific commands	395
Array specific commands	399
Rank specific commands	407
Extent pool specific commands	418
Address group specific commands	427
Logical control unit specific commands	429
CKD logical volume specific commands	438
Logical subsystem specific commands	457
Fixed block logical volume specific commands	462
Volume group specific commands	482
Advanced operation commands	491
Copy Services commands	493
FlashCopy commands	493
Remote FlashCopy commands	511
Remote Mirror and Copy path commands	529
Remote Mirror and Copy pair commands	541
Global Mirror commands	573
Global Mirror session commands	587
Audit commands	595
offloadauditlog	595

Chapter 9. Command-line interface scenarios 599

Modifying fixed block volume groups	599
---	-----

Deleting data storage configurations	600
Deleting a fixed block data storage configuration	600
Deleting a count key data storage configuration	604
Processing remote FlashCopy (inband) transactions	607
Metro Mirror test scenario: failback operation from local to remote site	608
Allowed remote mirror and copy volume pair conversions	610
Chapter 10. Messages in the CLI and management console server	613
Accessibility	615
Notices	617
Terms and conditions.	618

Trademarks	619
Electronic emission notices	620
Federal Communications Commission (FCC) statement.	620
Industry Canada compliance statement.	620
European community compliance statement	620
Japanese Voluntary Control Council for Interference (VCCI) class A statement	621
Korean Ministry of Information and Communication (MIC) statement	622
Taiwan class A compliance statement	622
Taiwan Contact Information	622

Glossary	623
.	623

Index	653
------------------------	------------

About this guide

This publication introduces the IBM System Storage DS Command-Line Interface application and provides instructions for installing it.

This publication provides descriptions of the following components:

- Each command-line interface command.
- Each system message that appears in the command-line interface application.
- Each system message that appears in the management console application.

This publication is for the storage system administrator or the person who is responsible for installing and using the DS Command-Line Interface in your enterprise. It is written for a user who understands the concepts of a command-line interface application and writing scripts and who has a knowledge of the operating systems and the storage systems in your enterprise.

Notices and publication information

This section contains information about safety notices that are used in this guide, environmental notices for this product, publication information, and information about sending your comments to IBM.

Safety notices

Complete this task to find information about safety notices.

To find the translated text for a danger or caution notice:

1. Look for the identification number at the end of each danger notice or each caution notice. In the following examples, the numbers **1000** and **1001** are the identification numbers.

DANGER

A danger notice indicates the presence of a hazard that has the potential of causing death or serious personal injury.

1000

CAUTION:

A caution notice indicates the presence of a hazard that has the potential of causing moderate or minor personal injury.

1001

2. Find the number that matches in the *IBM System Storage™ Solutions Safety Notices for IBM Versatile Storage Server and IBM System Storage Enterprise Storage Server, GC26-7229*.

Environmental notices

This section identifies the environmental guidelines that pertain to this product.

Product recycling and disposal

This unit contains recyclable materials.

This unit must be recycled or discarded according to applicable local and national regulations. IBM® encourages owners of information technology (IT) equipment to responsibly recycle their equipment when it is no longer needed. IBM offers a variety of product return programs and services in several countries to assist equipment owners in recycling their IT products. Information on IBM product recycling offerings can be found on IBM's Internet site at <http://www.ibm.com/ibm/environment/products/prp.shtml>.



Notice: This mark applies only to countries within the European Union (EU) and Norway.

Appliances are labeled in accordance with European Directive 2002/96/EC concerning waste electrical and electronic equipment (WEEE). The Directive determines the framework for the return and recycling of used appliances as applicable throughout the European Union. This label is applied to various products to indicate that the product is not to be thrown away, but rather reclaimed upon end of life per this Directive.

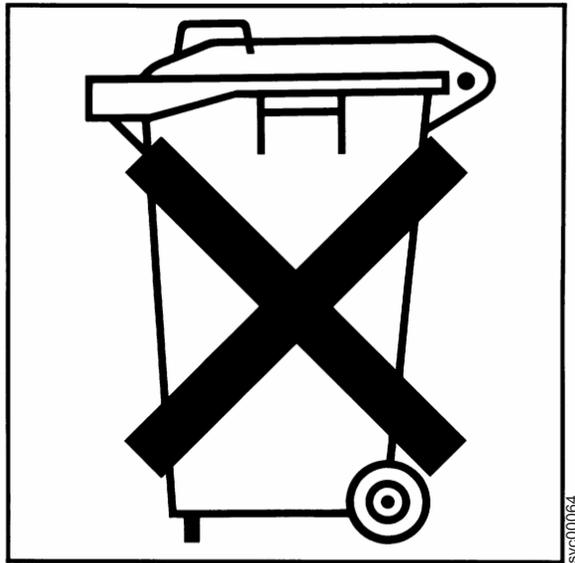
In accordance with the European WEEE Directive, electrical and electronic equipment (EEE) is to be collected separately and to be reused, recycled, or recovered at end of life. Users of EEE with the WEEE marking per Annex IV of the WEEE Directive, as shown above, must not dispose of end of life EEE as unsorted municipal waste, but use the collection framework available to customers for the return, recycling and recovery of WEEE. Customer participation is important to minimize any potential effects of EEE on the environment and human health due to the potential presence of hazardous substances in EEE. For proper collection and treatment, contact your local IBM representative.

Battery return program

This product may contain sealed lead acid, nickel cadmium, nickel metal hydride, lithium, or lithium ion battery. Consult your user manual or service manual for specific battery information. The battery must be recycled or disposed of properly. Recycling facilities may not be available in your area. For information on disposal of batteries outside the United States, go to <http://www.ibm.com/ibm/environment/products/batteryrecycle.shtml> or contact your local waste disposal facility.

In the United States, IBM has established a return process for reuse, recycling, or proper disposal of used IBM sealed lead acid, nickel cadmium, nickel metal hydride, and other battery packs from IBM Equipment. For information on proper disposal of these batteries, contact IBM at 1-800-426-4333. Please have the IBM part number listed on the battery available prior to your call.

In the Netherlands the following applies:



For Taiwan:



Please recycle batteries.

廢電池請回收

Conventions used in this guide

The following typefaces are used to show emphasis:

boldface

Text in **boldface** represents menu items and lowercase or mixed-case 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.

monospace

Text in monospace identifies the data or commands that you type, samples of command output, or examples of program code or messages from the system.

Related information

The tables in this section list and describe the following publications:

- The publications that make up the IBM® System Storage™ DS8000™ series library
- Other IBM publications that relate to the DS8000 series
- Non-IBM publications that relate to the DS8000 series

See “Ordering IBM publications” on page xvii for information about how to order publications in the IBM System Storage DS8000 series publication library. See “How to send your comments” on page xviii for information about how to send comments about the publications.

DS8000 series library

These customer publications make up the DS8000 series library.

Unless otherwise noted, these publications are available in Adobe portable document format (PDF) on a compact disc (CD) that comes with the storage unit. If you need additional copies of this CD, the order number is SK2T-8825. These publications are also available as PDF files by clicking on the **Documentation link** on the following Web site:

<http://www-1.ibm.com/servers/storage/support/disk>

See “Ordering IBM publications” on page xvii for information about ordering these and other IBM publications.

Title	Description	Order Number
<i>IBM System Storage DS: Command-Line Interface User's Guide</i>	This guide describes the commands that you can use from the command-line interface (CLI) for managing your DS8000 configuration and Copy Services relationships. The CLI application provides a set of commands that you can use to write customized scripts for a host system. The scripts initiate predefined tasks in a Copy Services server application. You can use the CLI commands to indirectly control Remote Mirror and Copy and FlashCopy® configuration tasks within a Copy Services server group.	SC26-7625 (See Note.)
<i>IBM System Storage DS8000: Host Systems Attachment Guide</i>	This guide provides guidelines for attaching the DS8000 to your host system and for migrating to fibre-channel attachment from a small computer system interface.	SC26-7628 (See Note.)
<i>IBM System Storage DS8000: Introduction and Planning Guide</i>	This guide introduces the DS8000 product and lists the features you can order. It also provides guidelines for planning the installation and configuration of the storage unit.	GC35-0495
<i>IBM System Storage Multipath Subsystem Device Driver User's Guide</i>	This publication describes how to use the IBM Subsystem Device Driver (SDD) on open-systems hosts to enhance performance and availability on the DS8000. SDD creates single devices (vpaths) that consolidate redundant paths for logical unit numbers. SDD permits applications to run without interruption when path errors occur. It balances the workload across paths, and it transparently integrates with applications.	SC30-4096
<i>IBM System Storage DS8000: User's Guide</i>	This guide provides instructions for setting up and operating the DS8000 and for analyzing problems.	SC26-7623 (See Note.)
<i>IBM System Storage DS Application Programming Interface Reference</i>	This publication provides reference information for the IBM System Storage DS application programming interface (API) and provides instructions for installing the Common Information Model Agent, which implements the API.	GC35-0493

Title	Description	Order Number
<i>IBM System Storage DS8000 Messages Reference</i>	This publication provides explanations of error, information, and warning messages that are issued from the DS8000 user interfaces.	GC26-7659
Note: No hardcopy book is produced for this publication. However, a PDF file is available from the following Web site: http://www-1.ibm.com/servers/storage/support/disk		

Other IBM publications

Other IBM publications contain additional information that is related to the DS product library.

The following list is divided into categories to help you find publications that are related to specific topics. Some of the publications are listed under more than one category. See “Ordering IBM publications” on page xvii for information about ordering these and other IBM publications.

Title	Description	Order Number
Data-copy services		
<i>z/OS DFSMS Advanced Copy Services</i>	This publication helps you understand and use IBM Advanced Copy Services functions. It describes three dynamic copy functions and several point-in-time copy functions. These functions provide backup and recovery of data if a disaster occurs to your data center. The dynamic copy functions are peer-to-peer remote copy, extended remote copy, and coupled extended remote copy. Collectively, these functions are known as remote copy. FlashCopy, SnapShot, and concurrent copy are the point-in-time copy functions.	SC35-0428
<i>IBM Enterprise Storage Server</i>	This publication, from the IBM International Technical Support Organization, introduces the Enterprise Storage Server and provides an understanding of its benefits. It also describes in detail the architecture, hardware, and functions, including the advanced copy functions, of the Enterprise Storage Server.	SG24-5465
<i>Implementing Copy Services On S/390</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on an Enterprise Storage Server that is attached to an S/390 or zSeries host system. Copy Services functions include peer-to-peer remote copy (PPRC), extended remote copy (XRC), FlashCopy®, and concurrent copy. This publication describes the functions, prerequisites, and corequisites and describes how to implement each function into your environment.	SG24-5680
<i>IBM TotalStorage ESS Implementing Copy Services in an Open Environment</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on UNIX, Windows NT®, Windows 2000, Sun Solaris, HP-UX, Tru64, OpenVMS, and iSeries host systems. The Copy Services functions that are described include peer-to-peer remote copy (PPRC) and FlashCopy. This publication describes the functions and shows you how to implement them into your environment. It also shows you how to implement these functions in a high-availability cluster multiprocessing environment.	SG24-5757
Fibre channel		

Title	Description	Order Number
<i>Fibre Channel Connection (FICON) I/O Interface: Physical Layer</i>	This publication provides information about the fibre-channel I/O interface. This book is also available as a PDF file from the following Web site: http://www.ibm.com/servers/resourcelink/	SA24-7172
<i>Fibre Transport Services (FTS): Physical and Configuration Planning Guide</i>	This publication provides information about fibre-optic and ESCON-trunking systems.	GA22-7234
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide</i>	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08.	SC26-7350
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide</i>	This guide describes the IBM SAN Fibre Channel Switch and the IBM TotalStorage ESS Specialist. It provides information about the commands and how to manage the switch with Telnet and the Simple Network Management Protocol.	SC26-7349
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide</i>	This publication describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S16. It is intended for trained service representatives and service providers.	SC26-7352
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide</i>	This guide introduces the IBM SAN Fibre Channel Switch 2109 Model S16 and tells you how to manage and monitor the switch using zoning and how to manage the switch remotely.	SC26-7351
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the Enterprise Storage Server. It provides you with a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment.	SG24-6113
Open-systems hosts		
<i>ESS Solutions for Open Systems Storage: Compaq AlphaServer, HP, and Sun</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure the Enterprise Storage Server when you attach Compaq AlphaServer (running Tru64 UNIX), HP, and Sun hosts. This book does not cover Compaq AlphaServer that is running the OpenVMS operating system. This book also focuses on the settings that are required to give optimal performance and on the settings for device driver levels. This book is for the experienced UNIX professional who has a broad understanding of storage concepts.	SG24-6119
<i>IBM TotalStorage ESS Implementing Copy Services in an Open Environment</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on UNIX or Windows 2000 host systems. The Copy Services functions that are described include peer-to-peer remote copy and FlashCopy. This publication describes the functions and shows you how to implement them into your environment. It also shows you how to implement these functions in a high-availability cluster multiprocessing environment.	SG24-5757
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the Enterprise Storage Server. It gives you a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment.	SG24-6113

Title	Description	Order Number
S/390 and zSeries hosts		
<i>Device Support Facilities: User's Guide and Reference</i>	This publication describes the IBM Device Support Facilities (ICKDSF) product that are used with IBM direct access storage device (DASD) subsystems. ICKDSF is a program that you can use to perform functions that are needed for the installation, the use, and the maintenance of IBM DASD. You can also use it to perform service functions, error detection, and media maintenance.	GC35-0033
<i>z/OS Advanced Copy Services</i>	This publication helps you understand and use IBM Advanced Copy Services functions. It describes three dynamic copy functions and several point-in-time copy functions. These functions provide backup and recovery of data if a disaster occurs to your data center. The dynamic copy functions are peer-to-peer remote copy, extended remote copy, and coupled extended remote copy. Collectively, these functions are known as remote copy. FlashCopy, SnapShot, and concurrent copy are the point-in-time copy functions.	SC35-0428
<i>DFSMS/MVS V1: Remote Copy Guide and Reference</i>	This publication provides guidelines for using remote copy functions with S/390 and zSeries hosts.	SC35-0169
<i>Fibre Transport Services (FTS): Physical and Configuration Planning Guide</i>	This publication provides information about fibre-optic and ESCON-trunking systems.	GA22-7234
<i>Implementing ESS Copy Services on S/390</i>	This publication, from the IBM International Technical Support Organization, tells you how to install, customize, and configure Copy Services on an Enterprise Storage Server that is attached to an S/390 or zSeries host system. Copy Services functions include peer-to-peer remote copy, extended remote copy, FlashCopy, and concurrent copy. This publication describes the functions, prerequisites, and corequisites and describes how to implement each function into your environment.	SG24-5680
<i>ES/9000, ES/3090: IOCP User Guide Volume A04</i>	This publication describes the Input/Output Configuration Program that supports the Enterprise Systems Connection (ESCON) architecture. It describes how to define, install, and configure the channels or channel paths, control units, and I/O devices on the ES/9000 processors and the IBM ES/3090 Processor Complex.	GC38-0097
<i>IOCP User's Guide, IBM e(logo)server zSeries 800 and 900</i>	This publication describes the Input/Output Configuration Program that supports the zSeries 800 and 900 servers. This publication is available in PDF format by accessing ResourceLink at the following Web site: www.ibm.com/servers/resourcelink/	SB10-7029
<i>IOCP User's Guide, IBM e(logo)server zSeries</i>	This publication describes the Input/Output Configuration Program that supports the zSeries server. This publication is available in PDF format by accessing ResourceLink at the following Web site: www.ibm.com/servers/resourcelink/	SB10-7037
<i>S/390: Input/Output Configuration Program User's Guide and ESCON Channel-to-Channel Reference</i>	This publication describes the Input/Output Configuration Program that supports ESCON architecture and the ESCON multiple image facility.	GC38-0401
<i>IBM z/OS Hardware Configuration Definition User's Guide</i>	This guide provides conceptual and procedural information to help you use the z/OS Hardware Configuration Definition (HCD) application. It also explains: <ul style="list-style-type: none"> • How to migrate existing IOCP/MVSCP definitions • How to use HCD to dynamically activate a new configuration • How to resolve problems in conjunction with MVS/ESA HCD 	SC33-7988

Title	Description	Order Number
<i>OS/390: Hardware Configuration Definition User's Guide</i>	This guide provides detailed information about the input/output definition file and about how to configure parallel access volumes. This guide discusses how to use Hardware Configuration Definition for both OS/390® and z/OS V1R1.	SC28-1848
<i>OS/390 V2R10.0: MVS System Messages Volume 1 (ABA - ASA)</i>	This publication lists OS/390 MVS™ system messages ABA to ASA.	GC28-1784
<i>Using IBM 3390 Direct Access Storage in a VM Environment</i>	This publication provides device-specific information for the various models of the 3390 and describes methods you can use to manage storage efficiently using the VM operating system. It provides guidance on managing system performance, availability, and space through effective use of the direct access storage subsystem.	GG26-4575
<i>Using IBM 3390 Direct Access Storage in a VSE Environment</i>	This publication helps you use the 3390 in a VSE environment. It includes planning information for adding new 3390 units and instructions for installing devices, migrating data, and performing ongoing storage management activities.	GC26-4576
<i>Using IBM 3390 Direct Access Storage in an MVS Environment</i>	This publication helps you use the 3390 in an MVS environment. It includes device-specific information for the various models of the 3390 and illustrates techniques for more efficient storage management. It also offers guidance on managing system performance, availability, and space utilization through effective use of the direct access storage subsystem.	GC26-4574
<i>z/Architecture Principles of Operation</i>	This publication provides a detailed definition of the z/Architecture™. It is written as a reference for use primarily by assembler language programmers and describes each function at the level of detail needed to prepare an assembler language program that relies on a particular function. However, anyone concerned with the functional details of z/Architecture will find this publication useful.	SA22-7832
SAN		
<i>IBM OS/390 Hardware Configuration Definition User's Guide</i>	<p>This guide explains how to use the Hardware Configuration Data application to perform the following tasks:</p> <ul style="list-style-type: none"> • Define new hardware configurations • View and modify existing hardware configurations • Activate configurations • Query supported hardware • Maintain input/output definition files (IODFs) • Compare two IODFs or compare an IODF with an actual configuration • Print reports of configurations • Create graphical reports of a configuration • Migrate existing configuration data 	SC28-1848
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 Installation and Service Guide</i>	This guide describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S08.	SC26-7350
<i>IBM SAN Fibre Channel Switch: 2109 Model S08 User's Guide</i>	This guide describes the IBM SAN Fibre Channel Switch and the IBM TotalStorage ESS Specialist. It provides information about the commands and how to manage the switch with Telnet and the Simple Network Management Protocol (SNMP).	SC26-7349
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 Installation and Service Guide</i>	This publication describes how to install and maintain the IBM SAN Fibre Channel Switch 2109 Model S16. It is intended for trained service representatives and service providers.	SC26-7352

Title	Description	Order Number
<i>IBM SAN Fibre Channel Switch: 2109 Model S16 User's Guide</i>	This guide introduces the IBM SAN Fibre Channel Switch 2109 Model S16 and tells you how to manage and monitor the switch using zoning and how to manage the switch remotely.	SC26-7351
<i>Implementing Fibre Channel Attachment on the ESS</i>	This publication, from the IBM International Technical Support Organization, helps you install, tailor, and configure fibre-channel attachment of open-systems hosts to the Enterprise Storage Server. It provides you with a broad understanding of the procedures that are involved and describes the prerequisites and requirements. It also shows you how to implement fibre-channel attachment.	SG24-6113
Seascape family		
<i>IBM Enterprise Storage Server</i>	This publication, from the IBM International Technical Support Organization, introduces the Enterprise Storage Server and provides an understanding of its benefits. It also describes in detail the architecture, hardware, and functions, including the advanced copy functions, of the Enterprise Storage Server.	SG24-5465
<i>IBM Enterprise Storage Server Performance Monitoring and Tuning Guide</i>	This guide, from the IBM International Technical Support Organization, provides guidance on the best way to configure, monitor, and manage your Enterprise Storage Server to ensure optimum performance.	SG24-5656
<i>IBM Versatile Storage Server: Introduction and Planning Guide</i>	This publication introduces the IBM Versatile Storage Server™ and lists the features you can order. It also provides planning information for both 2105 Models B09 and 100.	GC26-7223
<i>Implementing the IBM Enterprise Storage Server in Your Environment</i>	This publication, from the IBM International Technical Support Organization, can help you install, tailor, and configure the Enterprise Storage Server in your environment.	SG24-5420
Storage management		
<i>Device Support Facilities: User's Guide and Reference</i>	This publication describes the IBM Device Support Facilities (ICKDSF) product used with IBM direct access storage device (DASD) subsystems. ICKDSF is a program that you can use to perform functions that are needed for the installation, the use, and the maintenance of IBM DASD. You can also use it to perform service functions, error detection, and media maintenance.	GC35-0033
<i>IBM TotalStorage Solutions Handbook</i>	This handbook, from the IBM International Technical Support Organization, helps you understand what makes up enterprise storage management. The concepts include the key technologies that you must know and the IBM subsystems, software, and solutions that are available today. It also provides guidelines for implementing various enterprise storage administration tasks so that you can establish your own enterprise storage management environment.	SG24-5250

Ordering IBM publications

You can order copies of IBM publications using the IBM publications center.

IBM publications center

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

The IBM publications center offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download free of charge. You can also order publications. The publications center displays prices in your local currency. You can access the IBM publications center through the following Web site:

<http://www.elink.ibm.com/public/applications/publications/cgi-bin/pbi.cgi>

Note: Open the Web site in a new browser window by right clicking on the link and selecting "Open in New Window."

Web sites

The following Web sites provide information about the IBM System Storage DS8000 series and other IBM storage products.

Type of Storage Information	Web Site
Concurrent Copy for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
Copy Services command-line interface (CLI)	http://www-1.ibm.com/servers/storage/support/software/cscli/
DS8000 Information Center	http://publib.boulder.ibm.com/infocenter/ds8000ic/index.jsp
DS8000 series publications	http://www-1.ibm.com/servers/storage/support/disk Click Documentation .
FlashCopy for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
Host system models, operating systems, and adapters that the storage unit supports	http://www.ibm.com/servers/storage/disk/ds8000/ Click Interoperability matrix .
IBM Disk Storage Feature Activation (DSFA)	http://www.ibm.com/storage/dsfa
IBM storage products	http://www.storage.ibm.com/
IBM System Storage DS8000 series	http://www-1.ibm.com/servers/storage/disk/ds8000
IBM version of the Java (JRE) that is often required for IBM products	http://www-106.ibm.com/developerworks/java/jdk/
Multiple Device Manager (MDM)	http://www.ibm.com/servers/storage/support/ Click Storage Virtualization .
Remote Mirror and Copy (formerly PPRC) for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/
SAN fibre channel switches	http://www.ibm.com/storage/fcswitch/
Storage Area Network Gateway and Router	http://www-1.ibm.com/servers/storage/support/san/
Subsystem Device Driver (SDD)	http://www-03.ibm.com/servers/storage/support/software/sdd
Technical notes and product tips	http://www.ibm.com/servers/storage/support/disk/ds8100/ Click Technical notes on the Troubleshooting tab.
z/OS Global Mirror (formerly XRC) for S/390 and zSeries host systems	http://www.storage.ibm.com/software/sms/sdm/

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this information or any other DS8000 series documentation, you can submit them in the following ways:

- e-mail

Submit your comments electronically to the following e-mail address:

starpubs@us.ibm.com

Be sure to include the name and order number of the book and, if applicable, the specific location of the text you are commenting on, such as a page number or table number.

- Mail

Fill out the Readers' Comments form (RCF) at the back of this book. Return it by mail or give it to an IBM representative. If the RCF has been removed, you can address your comments to:

International Business Machines Corporation
RCF Processing Department
Department 61C
9032 South Rita Road
TUCSON AZ 85775-4401

Summary of changes for SC26-7916-01 IBM System Storage DS8000 Command-Line Interface User's Guide

This document contains terminology, maintenance, and editorial changes. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

New Information

The following commands and topics were added or enhanced in this release:

- The "Activating your machine and feature licenses using the DS CLI" in Chapter 3 was upgraded for the use of some additional activation codes.
- The "Creating a default configuration setup with a profile file" in Chapter 3 was upgraded to include a typical edit process for your profile configuration file.
- "Using forced failover and failback during an unplanned Metro/Global Mirror outage" was added to Chapter 7.
- "Using forced failover and failback during a planned Metro/Global Mirror outage" was added to Chapter 7.
- **lskey** command - Enhanced for the addition of some new activation codes.
- **failoverpprc** command - Enhanced with the addition of the **-force** parameter.
- **failbackpprc** command - Enhanced with the addition of the **-force** parameter.
- **showsu** command - Enhanced the reports.
- **showsi** command - Enhanced the reports.
- **chsu** command - Removed all unsupported functions.
- **chsi** command - Removed all unsupported functions.
- **setmpw** command - New command used when DS8000 uses the IBM TotalStorage Productivity Center Replication Manager function.
- **lspprc** command - This command was updated to include descriptions for the reason codes that might be reported.

Chapter 1. Introduction to IBM System Storage DS8000 series

IBM System Storage DS8000 is a high-performance, high-capacity series of disk storage that is designed to support continuous operations. DS8000 series models (machine type 2107/242x) use the IBM POWER5™ server technology that is integrated with the IBM Virtualization Engine™ technology. DS8000 series models consist of a storage unit and one or two management consoles, two being the recommended configuration. You can attach a maximum of four DS8000s to each management console. The graphical user interface (GUI) or the command-line interface (CLI) allows you to logically partition storage and use the built-in Copy Services functions. For high-availability, hardware components are redundant.

You can read the following information to familiarize yourself with the DS8000 series:

- DS8000 series machine types
- DS8000 series models, including a model comparison chart
- DS8000 series performance features
- DS8000 series interfaces
- DS8000 series hardware specifics
- Supported systems for open systems, S/390®, and zSeries® hosts
- Data management elements
- Copy Services
- Data migration

You can also view the DS8000 e-learning overview from the DS8000 Storage Manager Welcome page or at the following Web site:

<http://www-1.ibm.com/support/docview.wss?rs=1113&uid=ssg1S7001199>

The e-learning overview provides an animated presentation about installation and configuration, service and support, and management tasks.

Overview of the DS8000 series models

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

DS8100

The DS8100 (Models 921 and 931) features a dual two-way processor complex and support for one expansion frame.

DS8300

The DS8300 (Models 922, 932, 9A2, and 9B2) features a dual four-way processor complex and support for one or two expansion frames. Models 9A2 and 9B2 support two IBM System Storage System logical partitions (LPARs) in one storage unit.

DS8000 expansion models (Models 92E, 9AE) can be attached to a base model to expand its capabilities:

- 921 and 931 (base) can attach to 92E (expansion)
- 922 and 932 (base) can attach to 92E (expansion)

9A2 and 9B2 (base) can attach to 9AE (expansion)

Note: Models 921, 922, and 9A2 apply only to hardware machine type 2107. These models and all other 2107 hardware models have been withdrawn from marketing. Although you can no longer order these hardware models directly from IBM, you can order add-on features or licensed functions for your existing 2107 hardware models. You can obtain the hardware models on an as-available basis through IBM Business Partners.

Figure 1 provides a high-level view of the components of a base model (Models 921, 922, 931, 932, 9A2, 9B2).

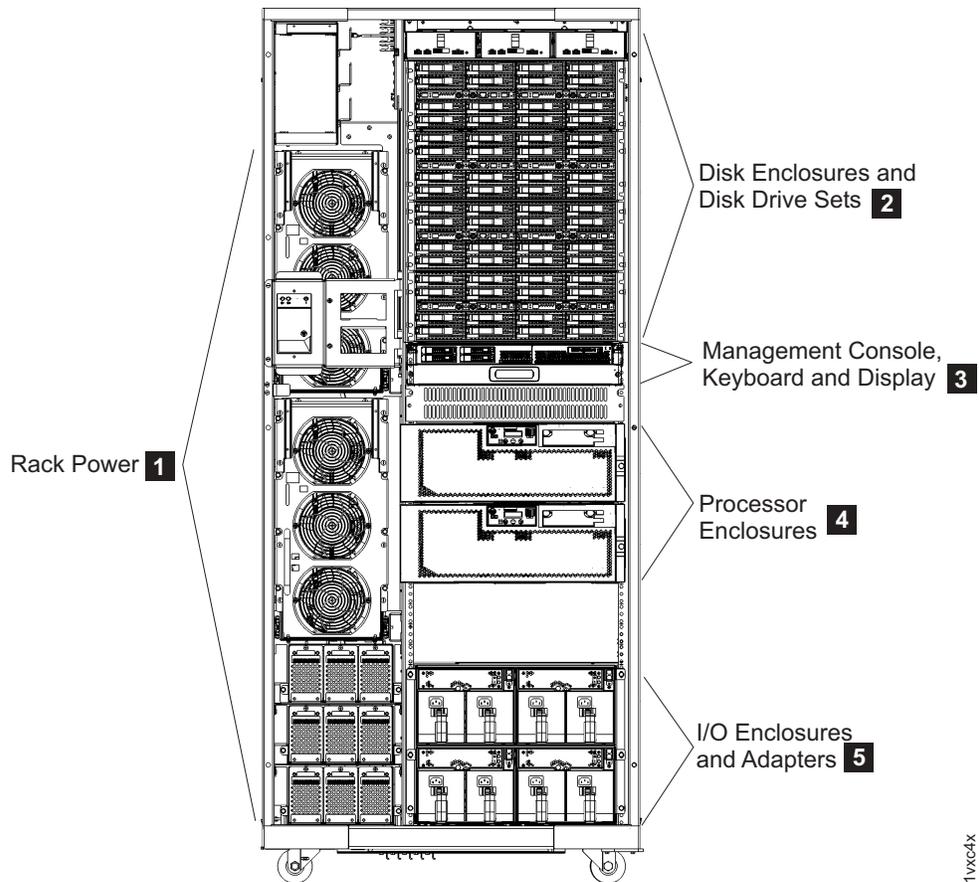


Figure 1. A base model (with front cover off) and its main components

The following notes provide additional information about the labeled components on Figure 1:

1. The *rack power* area of the base models provides redundant power supplies (two primary power supplies), power control cards, and backup battery assemblies to help protect data in the event of a loss of external power. Models 921 and 931 contain two batteries. Models 922, 932, 9A2, and 9B2 contain three batteries to support the 4-way processors.
2. All base models allow up to eight *disk enclosures*, which contain the disk drives. In a maximum configuration, each base model can hold up to 128 disk drives.

3. All base models contain one *management console*, a keyboard, and display. The management console is the focal point for configuration, copy services management, and maintenance activities.
4. All base models contain two *processor enclosures*. The Model 921 and 931 processor enclosures have 2-way processors. Processor enclosures on Models 922, 932, 9A2, and 9B2 have 4-way processors.
5. All base models contain *I/O enclosures and adapters*. The I/O enclosures hold the adapters and provide connectivity between the adapters and the processors. Both device adapters and host adapters are installed in the I/O enclosure.

Figure 2 provides a high-level view of the components of an expansion model (Models 92E, 9AE).

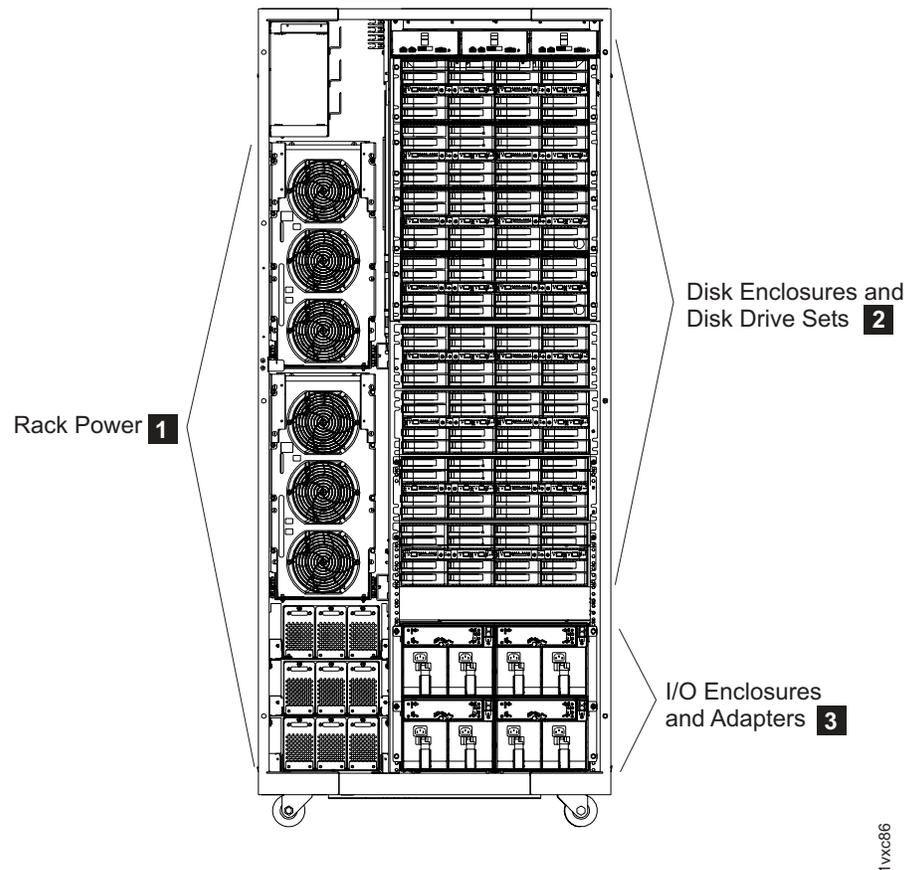


Figure 2. An expansion model (with front cover off) and its main components

The following notes provide additional information about the labeled components on Figure 2:

1. The *rack power* area of each expansion model provides redundant power supplies (two primary power supplies) and power control cards. If the expansion unit contains I/O enclosures or the extended power line disturbance (PLD) feature, two or three backup battery assemblies are also installed to help protect data in the event of a loss of external power.
2. All expansion models allow up to 16 *disk enclosures*, which contain the disk drives. In a maximum configuration, each expansion model can hold 256 disk drives.

- Expansion models can contain *I/O enclosures* and *adapters* if they are the first expansion units that are attached to a Model 922, 932, 9A2, or 9B2. The second expansion model in a 922, 932, 9A2, or 9B2 configuration cannot have I/O enclosures and adapters, nor can any expansion unit that is attached to a Model 921 or 931. If the expansion unit contains I/O enclosures, the enclosures provide connectivity between the adapters and the processors. The adapters contained in the I/O enclosures can be either device or host adapters.

DS8100 (Model 921 and Model 931)

The IBM System Storage DS8100, which is Model 921 and Model 931, offers many features.

These features include:

- Dual 2-way processor complex
- Up to 128 disk drives, for a maximum capacity of 64 TB
- Up to 128 GB of processor memory (cache)
- Up to 16 fibre channel/FICON or ESCON[®] host adapters

The DS8100 model can support one expansion frame. With one expansion frame, you can expand the capacity of the Model 921 or Model 931 as follows:

- Up to 384 disk drives, for a maximum capacity of 192 TB

Note: IBM service representatives can upgrade a Model 921 or Model 931 in the field when you order a model conversion to a Model 922, Model 932, Model 9A2, or Model 9B2. There are no available model conversions from Models 92x/9Ax to Models 93x/9Bx.

Figure 3 on page 5 shows the maximum configuration of a Model 921 or Model 931, which is the 921 or the 931 base model plus one 92E expansion model.

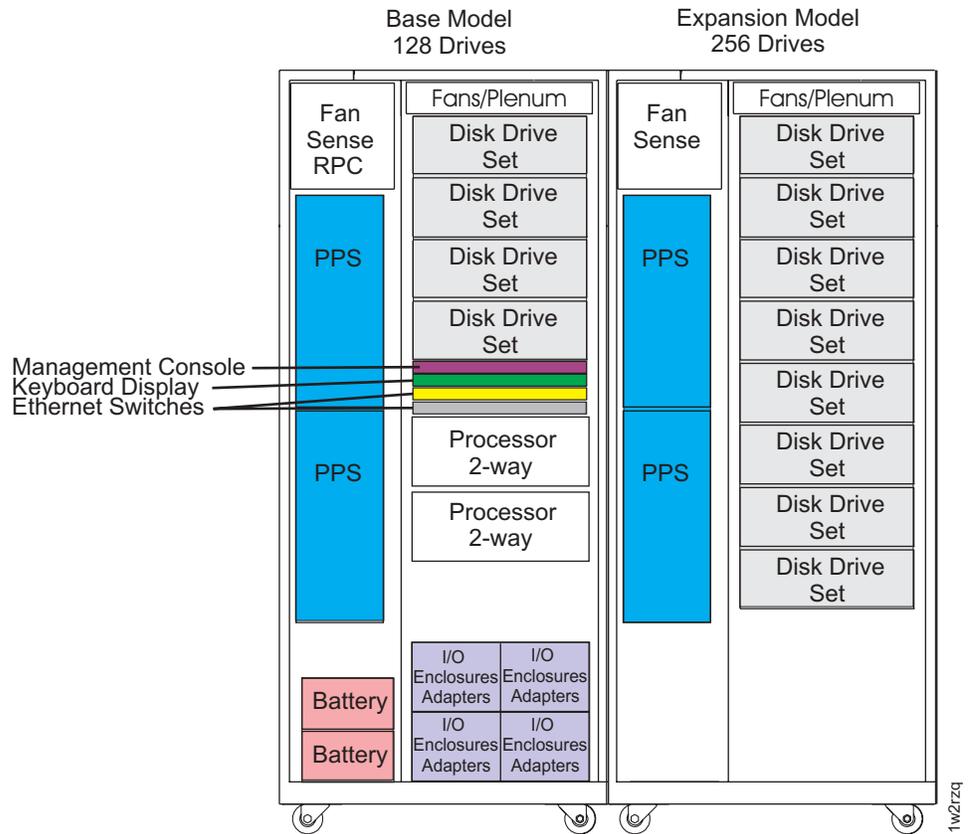


Figure 3. Maximum configuration for Model 921 or Model 931

DS8300 (Models 922, 932, 9A2, and 9B2)

IBM System Storage DS8300 models (Models 922, 932, 9A2, and 9B2) offer higher performance and capacity than the DS8100. Models 9A2 and 9B2 also enable you to create two storage system LPARs (images) within the same storage unit.

All DS8300 models offer the following features:

- Dual 4-way processor complex
- Up to 128 disk drives, for a maximum capacity of 64 TB
- Up to 256 GB of processor memory (cache)
- Up to 16 fibre-channel/FICON or ESCON host adapters

The DS8300 models can support either one or two expansion frames. With expansion frames, you can expand the Models 922, 932, 9A2, and 9B2 as follows:

- With one expansion frame, you can support the following expanded capacity and number of adapters:
 - Up to 384 disk drives, for a maximum capacity of 192 TB
 - Up to 32 fibre-channel/FICON or ESCON host adapters
- With two expansion frames, you can support the following expanded capacity:
 - Up to 640 disk drives, for a maximum capacity of 320 TB

Figure 4 shows the maximum configuration of a Model 922, 932, 9A2, or 9B2. All of these models can attach up to two expansion models. Models 922 and 932 can attach to the 92E expansion models. Models 9A2 and 9B2 can attach to the 9AE expansion models.

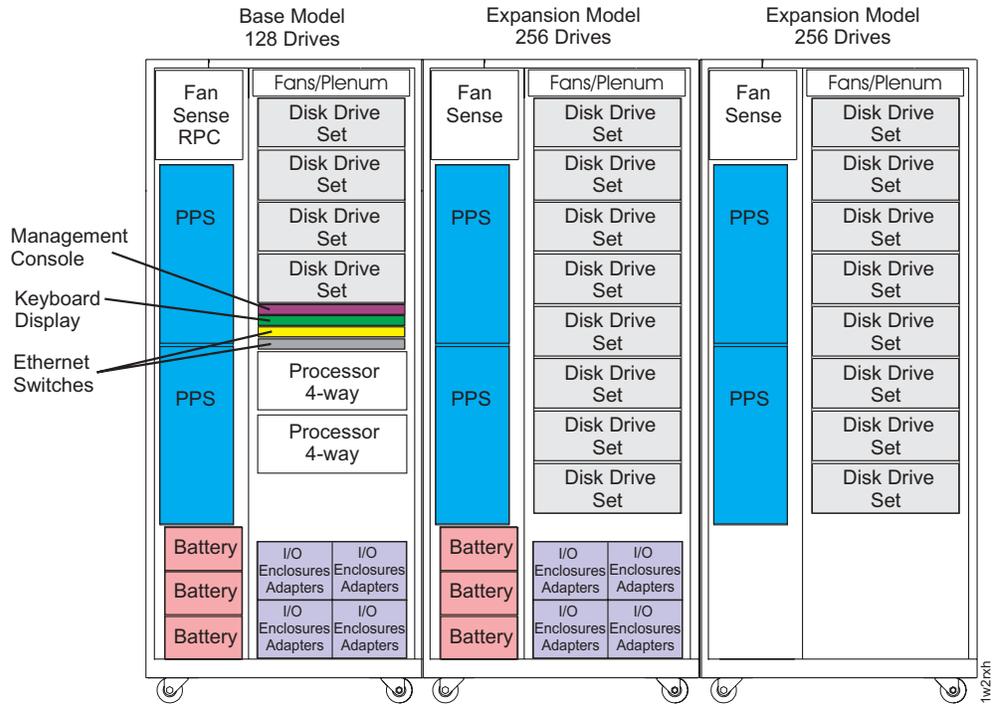


Figure 4. Maximum configuration for Model 922, 932, 9A2, or 9B2

Model comparison

DS8000 models vary on the processor types that they use, their disk capacity, and other supported features.

Table 1 compares the various supported features that are available on different DS8000 model configurations.

Table 1. DS8000 model comparison chart

Base model	Images (logical partitions)	Expansion models	Processor type	Disk drives	Processor memory	Host adapters
Model 921/931	1 image	None	2-way	Up to 128 disks (maximum capacity of 64 TB)	Up to 128 GB	Up to 16 host adapters
		With 1 expansion model (92E)	2-way	Up to 384 disks (maximum capacity of 192 TB)	Up to 128 GB	Up to 16 host adapters

Table 1. DS8000 model comparison chart (continued)

Base model	Images (logical partitions)	Expansion models	Processor type	Disk drives	Processor memory	Host adapters
Model 922/932	1 image	None	4-way	Up to 128 disks (maximum capacity of 64 TB)	Up to 256 GB	Up to 16 host adapters
		With 1 expansion model (92E)	4-way	Up to 384 disks (maximum capacity of 192 TB)	Up to 256 GB	Up to 32 host adapters
		With 2 expansion models (92E)	4-way	Up to 640 disks (maximum capacity of 320 TB)	Up to 256 GB	Up to 32 host adapters
Model 9A2/9B2	2 images	None	4-way	Up to 128 disks (maximum capacity of 64 TB)	Up to 256 GB	Up to 16 host adapters
		With 1 expansion model (9AE)	4-way	Up to 384 disks (maximum capacity of 192 TB)	Up to 256 GB	Up to 32 host adapters
		With 2 expansion models (9AE)	4-way	Up to 640 disks (maximum capacity of 320 TB)	Up to 256 GB	Up to 32 host adapters

Note: IBM service representatives can perform the conversion in the field when you order one of the following model conversions:

- Model 921 to 922
- Model 921 to 9A2
- Model 931 to 932
- Model 931 to 9B2
- Model 9A2 to 922
- Model 922 to 9A2
- Model 932 to 9B2
- Model 9B2 to 932
- Model 92E to 9AE
- Model 9AE to 92E

IBM System Storage Management Console

The IBM System Storage management console is the focal point for configuration, copy services management, and maintenance activities.

The management console is a dedicated workstation that is physically located (installed) inside your DS8100 and DS8300 and can automatically monitor the state of your system, notifying you and IBM when service is required. The management console also can be connected to your network to enable centralized management of your system using the IBM System Storage DS command-line interface or the storage management software that uses the IBM System Storage DS open API.

The first management console in a storage complex is always internal to the 2107/242x. To provide continuous availability of customer access to the management console functions, having a second management console is recommended. This second management console can be provided in two ways:

- **External** (outside the 2107/242x) - This console is generally installed in the customer-provided rack. It uses the same hardware as the internal management console.
- **Internal** (a second management console that resides within a second 2107/242x) - This console can be cabled together with the primary internal management console of the first 2107/242x.

The management console is the workstation the DS Storage Manager software is installed on. This console is used to configure and manage the DS8000 series system.

DS8000 Performance features

The DS8000 series is designed to provide you with high-performance, connectivity, and reliability allowing your workload to be easily consolidated into a single storage subsystem.

The following list provides an overview of some of the features that are associated with the DS8000 series:

POWER5+ processor technology

The DS8000 features IBM POWER5+ server technology. Compared with the current IBM POWER5 processor, the POWER5+ processor allows up to a 15% performance improvement in I/O operations per second in transaction processing workload environments. The POWER5+ processor is included only in the 93x/9Bx models, and it is an optional feature in the 92x/9Ax models.

Industry standard fibre-channel disk drives

The DS8000 series offers a selection of fibre-channel disk drives, including 300 GB fibre channel drives and 500 GB fibre-channel ATA (FATA) disk drives that allow a DS8100 to scale up to a capacity of 192 TB and a DS8300 to scale up to a capacity of 320 TB. The 500 GB FATA disk drives are lower cost disk drives that are designed for applications that require disk drive duty cycles that are equal to or less than 20%. If you use the FATA disk drives for applications that require duty cycles greater than 20%, your application is likely to experience limited array throughput and reduced performance. The 500 GB FATA disk drives are designed to support applications that use fixed-content or reference-only data, applications that are used for data archival operations, or near-line applications that require large amounts of data. Intermix of the FATA disk drive sets and fibre-channel disk drive sets is not supported within a disk enclosure or at creation. However, intermix on the same DA loop of enclosures with different drive types is supported if the drives within each enclosure are the same.

Four-port fibre-channel/FICON adapters

The DS8000 series offers enhanced connectivity with four-port fibre-channel/FICON host adapters. Four GB fibre-channel/FICON host support offers improvement that is up to 50 percent of the MB/second throughput performance for a single port. This improved performance enables cost savings by reducing the number of required host ports. These 4 GB host adapters are offered in longwave and shortwave. They auto-negotiate to either 4 GB or 2 GB link speeds. This flexibility provides benefits that are offered by higher performance 4 GB SAN-based solutions, while maintaining compatibility with existing 2 GB infrastructures.

HOST attachments

The DS8100 offers up to 16 host adapters (64 FCP/FICON ports) and the DS8300 offers up to 32 host adapters (128 FCP/FICON ports), further increasing your ability to share your disk storage.

IBM Standby Capacity on Demand

The IBM Standby Capacity on Demand (Standby CoD) offering allows the installation of inactive disk drives that can be easily activated as business needs require. With this offering, up to four Standby CoD disk drive sets (64 disk drives) can be factory- or field-installed into your system. To activate, you logically configure the disk drives for use—a nondisruptive activity that does not require intervention from IBM. Upon activation of any portion of the Standby CoD disk drive set, you must place an order with IBM to initiate billing for the activated set. At that time, you can also order replacement Standby CoD disk drive sets. 500 GB, 7 200 rpm fibre-channel ATA (FATA) Standby Capacity on Demand (Standby CoD) disk drive sets are available as part of the Standby CoD offering.

Online Information Center

The online Information Center is an information database that provides you the opportunity to quickly familiarize yourself with the major aspects of the DS8000 and to easily recognize the topics for which you might require more information. It provides information regarding user assistance for tasks, concepts, reference, user scenarios, tutorials, and other types of user 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.

RAID implementation

RAID implementation improves data storage reliability and performance.

Redundant array of independent disks (RAID) is a method of configuring multiple disk drives in a storage subsystem for high availability and high performance. The collection of two or more disk drives presents the image of a single disk drive to the system. In the event of a single device failure, data can be read or regenerated from the other disk drives in the array.

With RAID implementation, the DS8000 series offers fault-tolerant data storage by storing the data in different places on multiple disk drive modules (DDMs). By placing data on multiple disks, input/output operations can overlap in a balanced way to improve the basic reliability and performance of the attached storage devices.

Physical capacity for the DS8000 series can be configured as RAID 5, RAID 10, or a combination of both. RAID 5 can offer excellent performance for most applications,

while RAID 10 can offer better performance for selected applications, in particular, high random write content applications in the open systems environment.

You can reconfigure RAID 5 disk groups as RAID 10 disk groups or vice versa.

RAID 5 overview

RAID 5 is a method of spreading volume data across multiple disk drives. The DS8000 series supports RAID 5 arrays.

RAID 5 increases performance by supporting concurrent accesses to the multiple DDMs within each logical volume. Data protection is provided by parity, which is stored throughout the drives in the array. If a drive fails, the data on that drive can be restored using all the other drives in the array along with the parity bits that were created when the data was stored.

RAID 10 overview

RAID 10 provides high availability by combining features of RAID 0 and RAID 1. The DS8000 series supports RAID 10 arrays.

RAID 0 increases performance by striping volume data across multiple disk drives. RAID 1 provides disk mirroring which duplicates data between two disk drives. By combining the features of RAID 0 and RAID 1, RAID 10 provides a second optimization for fault tolerance.

RAID 10 implementation provides data mirroring from one DDM to another DDM. RAID 10 stripes data across half of the disk drives in the RAID 10 configuration. The other half of the array mirrors the first set of disk drives. Access to data is preserved if one disk in each mirrored pair remains available. In some cases, RAID 10 offers faster data reads and writes than RAID 5 because it does not need to manage parity. However, with half of the DDMs in the group used for data and the other half used to mirror that data, RAID 10 disk groups have less capacity than RAID 5 disk groups.

DS8000 Interfaces

This section describes the following interfaces:

- IBM System Storage DS Storage Manager
- IBM System Storage DS Command-Line Interface (CLI)
- IBM System Storage DS Open application programming interface
- IBM TotalStorage Productivity Center for Disk
- IBM TotalStorage Productivity for Replication Manager

IBM System Storage DS Storage Manager

The IBM System Storage DS Storage Manager is an interface that is used to perform logical configurations and Copy Services management functions.

The DS Storage Manager software is installed on a computer (management console) that is provided with the DS8000.

You can install the DS Storage Manager as a GUI (graphical mode) or as an unattended (silent) mode installation for the supported operating systems. The unattended (silent) mode installation option allows you to run the installation program unattended. Use this method of installation to customize a response file

and issue a command from a command prompt window. The response file is a template on the IBM System Storage Storage Manager CD. The DS Storage Manager can be accessed from any location that has network access using a Web browser. It offers you the following choices that are available with the applicable license management code:

Simulated configuration

You install this component on your PC or the management console which provides the ability to create or modify logical configurations when your storage unit is disconnected from the network. After creating the configuration, you can save it and then apply it to a network attached storage unit at a later time.

This simulated configuration tool is installed on your server and can be used for the configuration of a DS8000 series at initial installation.

Real-time configuration

This component is preinstalled on your management console. It provides you with the ability to create logical configurations and use Copy Services features when your storage unit is attached to the network. This component provides you with real-time (online) configuration support.

Copy services

This component allows you to process copy services functions.

The DS command-line interface

The IBM System Storage DS command-line interface (CLI) enables open systems hosts to invoke and manage FlashCopy and Metro and Global Mirror functions through batch processes and scripts.

The DS CLI provides a full-function set of commands that allows you to check your storage unit configuration and to perform specific application functions, when necessary. Many of the values that are used or reported with the DS CLI commands are hexadecimal.

The following list highlights a few of the specific types of functions that you can perform with the DS command-line interface:

- Check and verify your storage unit configuration
- Check the current Copy Services configuration that is used by the storage unit
- Create new logical storage and Copy Services configuration settings
- Modify or delete logical storage and Copy Services configuration settings

Notes:

1. You can have a maximum of 256 clients that are connected to the same server and this permits 256 DS CLI sessions. However, you must not start more than 100 DS CLI sessions simultaneously. Starting more than 100 DS CLI sessions simultaneously can result in connection problems.
2. You cannot install the DS CLI on a Windows 64-bit operating system.

DS open application programming interface

The IBM System Storage DS Open Application Programming Interface (API) is a nonproprietary storage management client application that supports routine LUN management activities, such as LUN creation, mapping and masking, and the creation or deletion of RAID 5 and RAID 10 volume spaces. The DS Open API

supports these activities through the use of the Storage Management Initiative Specification (SMI-S), as defined by the Storage Networking Industry Association (SNIA).

The DS Open API helps integrate configuration management support into storage resource management (SRM) applications, which allow customers to benefit from existing SRM applications and infrastructures. The DS Open API also enables the automation of configuration management through customer-written applications. Either way, the DS Open API presents another option for managing storage units by complementing the use of the IBM System Storage DS Storage Manager web-based interface and the DS command-line interface.

You must implement the DS Open API through the IBM System Storage Common Information Model (CIM) agent, a middleware application that provides a CIM-compliant interface. The DS Open API uses the CIM technology to manage proprietary devices as open system devices through storage management applications. The DS Open API allows these storage management applications to communicate with a storage unit.

The DS Open API supports the IBM System Storage DS8000 and the IBM System Storage DS6000, and the IBM TotalStorage Enterprise Storage Server. It is available for the AIX, Linux, and Windows operating system environments and must be used on storage units that have fibre-channel ports.

IBM TotalStorage Productivity Center for Disk

IBM TotalStorage® Productivity Center for Disk (Disk Manager) is a storage management product that is designed for customers who want to reduce the complexities and costs of storage management, including the management of SAN-based storage.

The IBM TotalStorage Productivity Center allows you to manage storage area networks and heterogeneous storage from a single console. It lets you manage network storage components that are based on the Storage Management Initiative Specification, including

- IBM SAN Volume Controller
- IBM TotalStorage Enterprise Storage Server® (ESS)
- IBM TotalStorage disk storage units (DS4000, DS6000, and DS8000 series)

Disk Manager can do the following performance functions:

- Collect and store performance data and provides alerts
- Generate both standard and custom performance reports with the ability to graph data from those reports
- Help optimize storage allocation
- Provide volume contention analysis and recommend intelligent placement of new volumes

Through the use of data collection, threshold settings, and performance reports, performance can be monitored for the ESS, DS4000, DS6000, and DS8000, SAN Volume Controller, and other storage subsystems that support the SMI-S block service performance subprofile. The performance functions start with the data collection task, which captures performance statistics for the devices and stores the data in the database.

You can set thresholds for certain performance metrics depending on the type of device. Threshold checking is performed during data collection, and when performance is outside the specified bounds, alerts can be generated.

After performance data has been collected, you can configure Disk Manager to present graphical or text reports on the historical performance behavior of specified devices, or of specific internal components of these devices. The performance reports provide information on the performance metrics and can display performance data in graphical form.

For more information, visit the Multiple Device Manager Information Center from the IBM Publications Web site using the following URL:

<http://publib.boulder.ibm.com/infocenter/tsmdmv11/index.jsp>

IBM TotalStorage Productivity Center for Replication

IBM TotalStorage Productivity Center for Replication provides Copy Services functions such as the remote mirror and copy functions (Metro Mirror and Global Mirror) and the point-in-time function (FlashCopy).

TotalStorage Productivity Center for Replication provides a graphical interface that you can use for configuring and managing Copy Services functions across the DS8000 and Enterprise Storage Server (ESS) storage units. These data-copy services maintain consistent copies of data on source volumes that are managed by Replication Manager.

IBM TotalStorage Productivity Center for Replication V3.1 for FlashCopy, Metro Mirror, and Global Mirror support provided automation of administration and configuration of these services, operational control (starting, suspending, resuming), Copy Services tasks, and monitoring and managing of copy sessions.

TotalStorage Productivity Center for Replication is an option of the TotalStorage Productivity Center for Replication software program. If you are licensed for Copy Services functions and have the required Ethernet adapter pair (feature code 180x) that supports the communication between the storage units and Replication Manager, you can use TotalStorage Productivity Center for Replication to manage your data copy environment.

For more information, visit the Multiple Device Manager Information Center from the IBM Publications Web site using the following URL:

<http://publib.boulder.ibm.com/infocenter/tsmdmv11/index.jsp>

Host systems attachment overview

The DS8000 storage unit provides a variety of host attachments so that you can consolidate storage capacity and workloads for open-systems hosts, S/390 hosts, and eServer™ zSeries hosts.

The DS8100 Models 921 and 931 support a maximum of 16 host adapters and 4 device adapter pairs, and the DS8300 Models 922, 932, 9A2, and 9B2 support a maximum of 32 host adapters and 8 device adapter pairs.

You can configure the storage unit for any of the following system adapter types and protocols:

- Fibre-channel adapters, for support of Fibre Channel Protocol (FCP) and fibre connection (FICON[®]) protocol
- Enterprise Systems Connection Architecture[®] (ESCON) adapters

For fibre-channel attachments, you can establish zones. The zones must contain a single port that is attached to a system adapter with the desired number of ports that are attached to the storage unit. By establishing zones, you reduce the possibility of interactions between system adapters in switched configurations. You can establish the zones by using either of two zoning methods:

- Port number
- Worldwide port name

You can configure switch ports that are attached to the storage unit in more than one zone. This enables multiple host system adapters to share access to the storage unit fibre-channel ports. Shared access to a storage unit fibre-channel port might come from host platforms that support a combination of bus adapter types and the operating systems.

IBM System Storage Resiliency Family

IBM System Storage Resiliency Family is a set of products and features that can help you implement storage solutions to keep your business running 24 hours a day, 7 days a week. This set of products provide disaster recovery, data migration, and data duplication functions.

As a part of the IBM System Storage Resiliency family, the DS8000 supports Copy Service functions that contribute to the protection of your data. These functions are also supported on the IBM TotalStorage Enterprise Storage Server.

The following IBM System Storage Resiliency family functions are available as optional features:

- Point-in-time copy, which includes IBM System Storage FlashCopy
 - The FlashCopy function enables you to make point-in-time, full volume copies of data, so that the copies are immediately available for read or write access. For zSeries environments, you can also use FlashCopy function to perform data set level copies of your data.
- Remote mirror and copy, which includes the following functions:
 - IBM System Storage Metro Mirror (previously known as Synchronous PPRC)
 - Metro Mirror provides real-time mirroring of logical volumes between two DS8000 storage units that can be located up to 300 km from each other. It is a synchronous copy solution where write operations are completed on both copies (local and remote site) before they are considered to be done.
 - IBM System Storage Global Copy (previously known as PPRC Extended Distance)
 - Global Copy is a nonsynchronous long-distance copy function where incremental updates are sent from the local to the remote site on a periodic basis.
 - IBM System Storage Global Mirror (previously known as Asynchronous PPRC)
 - Global Mirror is a long-distance remote copy function across two sites using asynchronous technology. Global Mirror processing is designed to provide

support for virtually unlimited distance between the local and remote sites, with the distance typically limited only by the capabilities of the network and the channel extension technology.

- IBM System Storage Metro/Global Mirror (a combination of Metro Mirror and Global Mirror)

Metro/Global Mirror is a three-site remote copy solution, which uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

- Remote mirror and copy for z/OS[®], which includes IBM System Storage z/OS Global Mirror is available only for zSeries environments

The point-in-time and remote mirror and copy features are supported across various server environments such as zSeries, pSeries[®], iSeries[™] and servers from Sun and Hewlett-Packard.

You can manage these functions through a command-line interface called the DS CLI and a Web-based interface called the DS Storage Manager. The DS Storage Manager allows you to set up and manage the following types of data-copy functions from any point where network access is available:

Point-in-time copy (FlashCopy)

The FlashCopy function enables you to make point-in-time, full volume copies of data, with the copies immediately available for read or write access. For zSeries environments, you can also use the FlashCopy function to perform data set level copies of your data. You can use the copy with standard backup tools that are available in your environment to create backup copies on tape.

FlashCopy is an optional function. To use it, you must purchase one of the point-in-time 2244/239x function authorization features.

The FlashCopy function creates a copy of a source volume on the target volume. This copy is called a point-in-time copy. When you initiate a FlashCopy operation, a FlashCopy relationship is created between a source volume and target volume. A FlashCopy relationship is a *mapping* of the FlashCopy source volume and a FlashCopy target volume. This mapping allows a point-in-time copy of that source volume to be copied to the associated target volume. The FlashCopy relationship exists between this volume pair from the time that you initiate a FlashCopy operation until the storage unit copies all data from the source volume to the target volume or you delete the FlashCopy relationship, if it is a persistent FlashCopy.

One of the main benefits of the FlashCopy function is that the point-in-time copy is immediately available for creating a backup of production data. The target volume is available for read and write processing so it can be used for testing or backup purposes. Data is physically copied from the source volume to the target volume using a background process. (A FlashCopy operation without a background copy is also possible, which allows only data that is modified on the source to be copied to the target volume.) The amount of time that it takes to complete the background copy depends on the following criteria:

- The amount of data being copied
- The number of background copy processes that are occurring
- The other activities that are occurring on the storage units

The FlashCopy function supports the following copy options:

Consistency groups

Creates a consistent point-in-time copy of multiple volumes, with negligible host impact. You can enable FlashCopy consistency groups from the DS CLI.

Change recording

Activates the change recording function on the volume pair that is participating in a FlashCopy relationship. This enables a subsequent refresh to the target volume.

Establish FlashCopy on existing Metro Mirror source

Allows you to establish a FlashCopy relationship where the target volume is also the source of an existing remote mirror and copy source volume. This enables you to create full or incremental point-in-time copies at a local site and then use remote mirroring commands to copy the data to the remote site.

Fast reverse

Reverses the FlashCopy relationship without waiting for the finish of the background copy of the previous FlashCopy. This option applies to the Global Mirror mode.

Inhibit writes to target

Ensures that write operations are inhibited on the target volume until a refresh FlashCopy operation is complete.

Multiple Relationship FlashCopy

Allows a source volume to have multiple (up to 12) target volumes at the same time.

Persistent FlashCopy

Allows the FlashCopy relationship to remain even after the FlashCopy operation completes. You must explicitly delete the relationship.

Refresh target volume

Provides the ability to refresh a FlashCopy relationship, without recopying all tracks from the source volume to the target volume.

Resynchronizing FlashCopy volume pairs

Provides the ability to update an initial point-in-time copy of a source volume without having to recopy your entire volume.

Reverse restore

Reverses the FlashCopy relationship and copies data from the target volume to the source volume.

Remote mirror and copy

The remote mirror and copy feature is a flexible data mirroring technology that allows replication between a source volume and a target volume on one or two disk storage units. You can also issue remote mirror and copy operations to a group of source volumes on one logical subsystem (LSS) and a group of target volumes on another LSS. (An LSS is a logical grouping of up to 256 logical volumes for which the volumes must have the same disk format, either count key data or fixed block.)

Remote mirror and copy is an optional feature that provides data backup and disaster recovery. To use it, you must purchase at least one of the remote mirror

and copy 2244/239x function authorization features. DS8000 storage units can participate in remote mirror and copy solutions with the DS6000 series, ESS Model 750, and ESS Model 800.

The remote mirror and copy feature provides synchronous (Metro Mirror) and asynchronous (Global Copy) data mirroring. The main difference is that the Global Copy feature can operate at very long distances, even continental distances, with minimal impact on applications. Distance is limited only by the network and channel extenders technology capabilities. The maximum supported distance for Metro Mirror is 300 km.

With Metro Mirror, application write performance is dependent on the available bandwidth. Global Copy allows you to better use your available bandwidth capacity, therefore allowing you to include more of your data to be protected.

The enhancement to Global Copy is Global Mirror, which uses Global Copy and the benefits of FlashCopy to form consistency groups. (A consistency group is a set of volumes that contain consistent and current data to provide a true data backup at a remote site.) Global Mirror uses a master storage unit (along with optional subordinate storage units) to internally, without external automation software, manage data consistency across volumes using consistency groups.

Consistency groups can also be created using the freeze and run functions of Metro Mirror. The freeze and run functions, when used with external automation software, provide data consistency for multiple Metro Mirror volume pairs.

The following sections describe the remote mirror and copy functions.

Synchronous mirroring (Metro Mirror)

Provides real-time mirroring of logical volumes (a source and a target) between two storage units that can be located up to 300 km from each other. With Metro Mirror copying, the source and target volumes can be on the same storage unit or on separate storage units. You can locate the storage unit at another site, some distance away.

Metro Mirror is a synchronous copy feature where write operations are completed on both copies (local and remote site) before they are considered to be complete. Synchronous mirroring means that a storage server constantly updates a secondary copy of a volume to match changes made to a source volume.

The advantage of synchronous mirroring is that there is minimal host impact for performing the copy. The disadvantage is that since the copy operation is synchronous, there can be an impact to application performance because the application I/O operation is not acknowledged as complete until the write to the target volume is also complete. The longer the distance between primary and secondary storage units, the greater this impact to application I/O, and therefore, application performance.

Asynchronous mirroring (Global Copy)

Copies data nonsynchronously and over longer distances than is possible with the Metro Mirror feature. When operating in Global Copy mode, the source volume sends a periodic, incremental copy of updated tracks to the target volume instead of a constant stream of updates. This causes less impact to application writes for source volumes and less demand for bandwidth resources, while allowing a more flexible use of the available bandwidth.

The updates are tracked and periodically copied to the target volumes. As a consequence, there is no guarantee that data is transferred in the same sequence that was applied to the source volume. To get a consistent copy of your data at your remote site, you must periodically switch from Global Copy to Metro Mirror mode, then either stop the application I/O or freeze data to the source volumes using a manual process with freeze and run commands. The freeze and run functions can be used with external automation software such as Geographically Dispersed Parallel Sysplex (GDPS), which is available for z/Series environments, to ensure data consistency to multiple Metro Mirror volume pairs in a specified logical subsystem.

Common options for Metro Mirror and Global Copy include the following modes:

Suspend and resume

If you schedule a planned outage to perform maintenance at your remote site, you can suspend Metro Mirror or Global Copy processing on specific volume pairs during the duration of the outage. During this time, data is no longer copied to the target volumes. Because the primary storage unit keeps track of all changed data on the source volume, you can resume operations at a later time to synchronize the data between the volumes.

Copy out-of-synchronous data

You can specify that only data that was updated on the source volume while the volume pair was suspended be copied to its associated target volume.

Copy an entire volume or not copy the volume

You can copy an entire source volume to its associated target volume to guarantee that the source and target volume contain the same data. When you establish volume pairs and elect not to copy a volume, a relationship is established between the volumes but no data is sent from the source volume to the target volume. In this case, it is assumed that the volumes contain exactly the same data and are consistent, so copying the entire volume is not necessary or required. Only new updates are copied from the source to target volumes.

Global Mirror

Provides a long-distance remote copy across two sites using asynchronous technology. Global Mirror processing is most often associated with disaster recovery or disaster recovery testing. However, it can also be used for everyday processing and data migration.

The Global Mirror function mirrors data between volume pairs of two storage units over greater distances without affecting overall performance. It also provides application-consistent data at a recovery (or remote) site in case of a disaster at the local site. By creating a set of remote volumes every few seconds, the data at the remote site is maintained to be a point-in-time consistent copy of the data at the local site.

Global Mirror operations periodically invoke point-in-time FlashCopy operations at the recovery site, at regular intervals, without disrupting the I/O to the source volume, thus giving a continuous, near up-to-date data backup. By grouping many volumes into a session, which is managed by the master storage unit, you can copy multiple volumes to the recovery site simultaneously while maintaining point-in-time consistency across

those volumes. (A session contains a group of source volumes that are mirrored asynchronously to provide a consistent copy of data at the remote site. Sessions are associated with Global Mirror relationships and are defined with an identifier [session ID] that is unique across the enterprise. The ID identifies the group of volumes in a session that are related and that can participate in the Global Mirror consistency group.)

Metro/Global Mirror

Provides a three-site, long distance disaster recovery replication that combines Metro Mirror with Global Mirror replication for both zSeries and open systems data. Metro/Global Mirror uses synchronous replication to mirror data between a local site and an intermediate site, and asynchronous replication to mirror data from an intermediate site to a remote site.

In a three-site, Metro/Global Mirror, should an outage occur, a backup site is maintained regardless of which one of the sites is lost. Suppose an outage occurs at the local site, Global Mirror continues to mirror updates between the intermediate and remote sites, maintaining the recovery capability at the remote site. If an outage occurs at the intermediate site, data at the local storage unit is not affected. If an outage occurs at the remote site, data at the local and intermediate sites is not affected. Applications continue to run normally in either case.

With the incremental resynchronization function enabled on a Metro/Global Mirror configuration, should the intermediate site be lost, the local and remote sites can be connected, and only a subset of changed data is copied between the volumes at the two sites. This reduces the amount of data that needs to be copied from the local site to the remote site and the time it takes to do the copy.

z/OS Global Mirror

In the event of workload peaks, which may temporarily overload the z/OS Global Mirror configuration's bandwidth, the enhanced z/OS Global Mirror function initiates a z/OS Global Mirror suspension preserving primary site application performance. If you are installing new high-performance z/OS Global Mirror primary storage subsystems, this function provides improved capacity and application performance during heavy write activity. This enhancement can also allow z/OS Global Mirror to be configured to tolerate longer periods of communication loss with the primary storage subsystems, allowing the z/OS Global Mirror to stay active despite transient channel path recovery events. In addition, this enhancement can provide fail-safe protection against application system impact related to unexpected data mover system events.

The z/OS Global Mirror function is an optional function. To use it, you must purchase the remote mirror for z/OS 2244/239x function authorization feature.

Disaster recovery using Copy Services

One of the main reasons for using Copy Services functions is to prepare for a possible disaster by backing up, copying, and mirroring your data both at the local (production) and remote sites.

Having a disaster recovery plan can ensure that critical data is recoverable at the time of a disaster. Because most disasters are unplanned, your disaster recovery plan must provide a way that allows you to recover your applications quickly, and

more importantly, to access your data. Consistent data to the same point-in-time across all storage units is vital before you can recover your data at a backup (normally your remote) site.

Most users use a combination of remote mirror and copy and point-in-time copy (FlashCopy) features to form a comprehensive enterprise solution for disaster recovery. In an event of a planned event or unplanned disaster, you can use failover and failback modes as part of your recovery solution. Failover and failback modes help to reduce the time that is required to synchronize remote mirror and copy volumes after you switch between the local (or production) and the intermediate or remote sites during planned and unplanned outages. Although failover transmits no data, it changes the status of a device, and the status of the secondary volume changes to a suspended primary volume. The Failback command transmits data and can go in either direction depending on which device the Failback command is issued to.

Recovery procedures that include failover and failback modes use remote mirror and copy functions, such as Metro Mirror, Global Copy, Global Mirror, Metro/Global Mirror, and FlashCopy.

Note: See the *IBM System Storage DS8000 Command-Line Interface User's Guide* for specific disaster recovery tasks.

Data consistency can be achieved using the following methods:

Manually using external software (without Global Mirror)

If you use Metro Mirror, Global Copy, and FlashCopy functions to create a consistent and restartable copy at your recovery site, you must do a manual and periodic suspend operation at your local site. This means using *freeze and run* commands together with external automated software and then using the FlashCopy function to make a consistent copy of your target volume for backup or recovery purposes. (Automation software is not provided with the storage unit; it must be supplied by the user.)

Note: Freezing of the data is done at the same point-in-time across all links and all storage units.

Automatically (with Global Mirror and FlashCopy)

If you use a two-site Global Mirror or a three-site Metro/Global Mirror configuration, the process to create a consistent and restartable copy at your intermediate or remote site is done using an automated process, with minimal or no interruption to your applications. Global Mirror operations automate the process of continually forming consistency groups. It combines Global Copy and FlashCopy operations to provide consistent data at the remote site. A master storage unit (along with subordinate storage units) internally manages data consistency using consistency groups within a Global Mirror configuration. Consistency groups can be created many times per hour to increase the currency of data that is captured in the consistency groups at the remote site.

Note: A consistency group is a collection of volumes (grouped in a session) across multiple storage units that are managed together in a session during the creation of consistent copies of data. The formation of these consistency groups is coordinated by the master storage unit, which sends commands over remote mirror and copy links to its subordinate storage units.

In a two-site Global Mirror configuration, if you have a disaster at your local site and have to start production at your remote site, you can use the consistent point-in-time data from the consistency group at your remote site to recover when the local site is operational.

In a three-site Metro/Global Mirror configuration, if you have a disaster at your local site and you must start production at either your intermediate or remote site, you can use the consistent point-in-time data from the consistency group at your remote site to recover when the local site is operational.

Comparison of licensed functions

A key decision that you must make in planning for a disaster is deciding which licensed functions to use to best suit your environment.

Table 2 provides a brief summary of the characteristics of the Copy Services features that are available for the storage unit.

Table 2. Comparison of licensed functions

Licensed function	Description	Advantages	Considerations
Metro/Global Mirror	Three-site, long distance disaster recovery replication	A backup site is maintained regardless of which one of the sites is lost.	Recovery point objective (RPO) might grow if bandwidth capability is exceeded.
Metro Mirror	Synchronous data copy at a distance	No data loss, rapid recovery time for distances up to 300 km.	Slight performance impact.
Global Copy	Continuous copy without data consistency	Nearly unlimited distance, suitable for data migration, only limited by network and channel extenders capabilities.	Copy is normally fuzzy but can be made consistent through synchronization.
Global Mirror	Asynchronous copy	Nearly unlimited distance, scalable, and low RPO. The RPO is the time needed to recover from a disaster; that is, the total system downtime.	RPO might grow when link bandwidth capability is exceeded.
z/OS Global Mirror	Asynchronous copy controlled by z/OS host software	Nearly unlimited distance, highly scalable, and very low RPO.	Additional host server hardware and software is required. The RPO might grow if bandwidth capability is exceeded or host performance might be impacted.

Logical configuration overview

It is important to understand IBM terminology for storage concepts and the storage hierarchy.

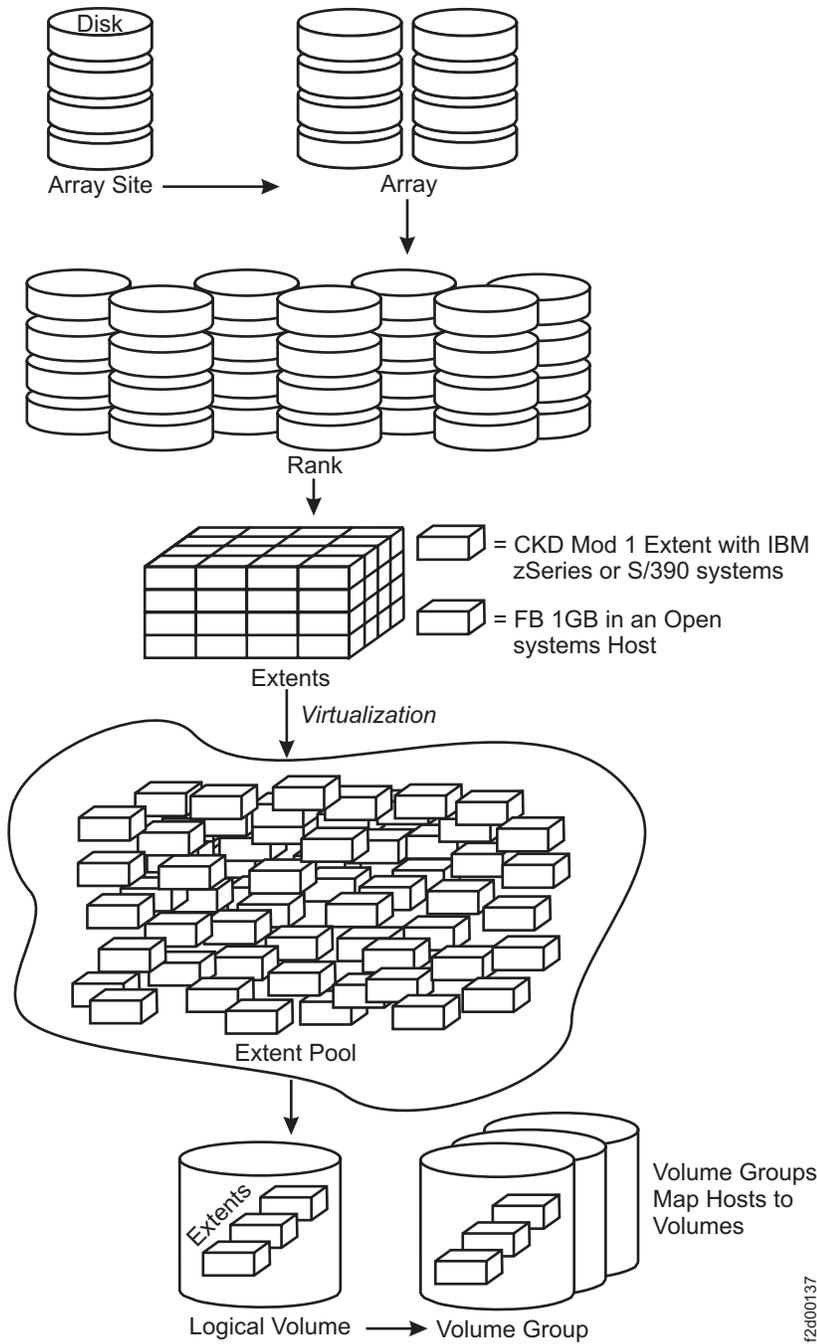
In the storage hierarchy, you begin with a disk. Logical groupings of four disks form an array site. Logical groupings of one or two array sites form an array. When you define your array storage type as CKD or fixed block, you create a rank. A rank is divided into a number of fixed-size extents. If you work with an open-systems host, an extent is 1 GB. If you work with IBM eServer zSeries or S/390 systems, an extent is the size of an IBM 3390 Mod 1 disk drive.

After you create ranks, your physical storage can be considered virtualized. Virtualization dissociates your physical storage configuration from your logical configuration, so that volume sizes are no longer constrained by the physical size of your arrays.

The available space on each rank is divided into extents. The extents are the building blocks of the logical volumes. An extent is striped across all disks of an array.

Extents of the same storage type are grouped together to form an extent pool. Multiple extent pools can create storage classes that provide greater flexibility in storage allocation through a combination of RAID types, DDM size, DDM speed, and DDM technology. This allows a differentiation of logical volumes by assigning them to the appropriate extent pool for the desired characteristics. Different extent sizes for the same device type (for example, count-key-data or fixed block) can be supported on the same storage unit, but these different extent types must be in different extent pools.

A logical volume is composed of one or more extents. A volume group specifies a set of logical volumes. By identifying different volume groups for different uses (for example, SCSI target, FICON/ESCON control unit, Remote Mirror and Copy secondaries, FlashCopy targets, and Copy Services), access to the set of logical volumes that are identified by the volume group can be controlled. Volume groups map hosts to volumes.



f2d00137

Figure 5. Logical configuration sequence

The storage management software can be used in Real-time or Simulated modes. When you are connected to storage devices over your network, you can use the Real-time Manager to manage your hardware or configure your storage. When you want to work offline, you can use the Simulated Manager to create or import a configuration. This lets you export a configuration that you have already created to a system when it reaches your site.

Chapter 2. Upgrading your system to use DS CLI

You can upgrade your system to use DS CLI by following the remove and install procedures put in place for DS CLI. However, there are some things that you must consider before you make this upgrade.

It is recommended that you upgrade the level of code and then the DS Storage Manager before upgrading your system to use DS CLI.

The DS CLI is designed to support the following features that exist on the IBM TotalStorage Enterprise Storage Server® (ESS) Models 750 and 800:

- A Copy Services domain, configured as part of the IBM TotalStorage Enterprise Storage Server (ESS) Models 750 and 800.
- FlashCopy Version 2 and PPRC Version 2 licenses that are available on the ESS.

However, the DS CLI is not designed to support the CLI scripts that you have written for these features without modification.

As part of your upgrade preparation, consider the following tips and questions:

- An upgrade requires that you uninstall the DS CLI and that you reinstall the upgraded DS CLI. This is the most certain way to ensure that you receive a clean installation. However, this uninstall and reinstall process can be a concern where you have customized the system profile file. A reinstallation will most likely overwrite your current system profile file. If you want to keep your current system profile file, do the following actions:
 - Make a copy of your current system profile file and save it in a convenient place.
 - Copy the saved system profile file into the DS CLI installation directory after the installation has completed. This allows you to overwrite the system profile file that was installed with the upgraded DS CLI.

- Do you plan to add the DS8000 to your network?

If you do, you must install the DS CLI application. This application is used for the configuration, performance, and Copy Services functions. Also, only the DS CLI can be used to perform Copy Services functions between the ESS and the DS models.

- Do you plan to keep the DS8000 separate from your network?

If you do, you cannot communicate with the ESSs in your network, even for Copy Services functions. You can install the DS CLI on your server and it supports the DS8000, but new scripts must be written that support the ESSs.

Chapter 3. Installing and removing the CLI

On most systems you can install and remove the DS command-line interface (CLI) using silent mode, console mode, or by using a GUI application.

Supported operating systems for the DS command-line interface

The DS command-line interface (CLI) can be installed on these operating systems.

- AIX 5.1, 5.2, 5.3
- HP-UX 11.0, 11i, v1, v2
- HP Tru64 version 5.1, 5.1A
- Linux (RedHat 3.0 Advanced Server (AS) and Enterprise Server (ES))
- SUSE Linux SLES 8, SLES 9, SUSE 8, SUSE 9
- Novell Netware 6.5
- System i system i5/OS 5.3
- OpenVMS 7.3-1 (or newer)
- Sun Solaris 7, 8, 9
- Windows 2000, Windows Datacenter, Windows 2003, and Windows XP

Note: The DS CLI cannot be installed on a Windows 64-bit operating system.

DS CLI operational limitations

There are operational limitations 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 logical control units (LCUs), which are the CKD volume equivalent of a fixed block LSS.

- If you are using the Cisco MDS 9216 Multilayer Fabric Switch, you must not enable its write acceleration feature. The Remote Mirror and Copy commands might fail if the write acceleration feature is enabled.
- A maximum of 256 volumes for each logical subsystem can be defined.
- You must not start more than 100 DS CLI sessions simultaneously. Starting more than 100 DS CLI sessions simultaneously can result in connection problems.

Preparing for the DS CLI installation

The IBM System Storage DS CLI program enables open systems hosts to invoke and manage FlashCopy and Metro and Global Mirror functions through batch processes and scripts. You can choose to install the DS CLI interface using the unattended (silent) mode, console mode, or graphic mode. The mode of installation that you choose is often dependent on the configuration of your system.

The DS CLI can be installed on only one of the following supported operating systems:

Supported Host Systems	Installation File Name
IBM AIX (5.1, 5.2, 5.3)	setupaix.bin
Hewlett-Packard-UX (11.0, 11i, v1, v2)	setuphp1020.bin or setuphp11x.bin
Linux (Red Hat 3.0 Advanced Server [AS] and Enterprise Server [ES])	setuplinux.bin for xSeries setupgenericunix.sh for pSeries
SUSE Linux SLES 8, SLES 9, SUSE 8, SUSE 9	setuplinux.bin for xSeries setupgenericunix.sh for pSeries
Sun Solaris (7, 8, 9)	setupsolarisSparc.bin
HP Tru64 (5.1, 5.1A)	setupgenericunix.sh
Novell Netware 6.5	setupwin32.exe
System i system i5/OS 5.3	Add -os400 at a command line after the name of the installation file. For example, setupaix.bin -os400
Windows 2000, Windows Datacenter, Windows 2003, and Windows XP	setupwin32.exe
UNIX users that do not have an X display	setupos.bin -console
Note: Use the same installation file name for your host operating system as shown in the previous rows, but add the -console parameter after the installation file name. For example: setupaix.bin -console setuplinux.bin -console	Where <i>os</i> represents the name of your operating system as shown in the previous rows.

- The DS CLI cannot be installed on a Windows 64-bit system.
- You must have Java 1.4.1 or later installed on your machine. The installation program checks for this requirement during installation and does not install the DS CLI if you do not have Java 1.4.1 or later.

Note: The DS CLI installation CD-ROM contains Java 1.4.2, which you can install if your system is not equipped with this version of Java.

- You must use the ksh (Korn shell) or bash (Bourne again shell) shell if you are installing on a HP Tru64. Installshield does not support the sh shell. You must perform all HP Tru64 installations using the **setupgenericunix.sh** file that is located on the installation compact disc.
- AIX installation:
 - The LIBPATH environment variable can interfere with the installation of the DS CLI and can result in the display of the JVM-not-found error. To avoid

this interference, 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.

- Run the following commands to sequentially disable the LIBPATH environment variable, install the DS CLI, and restore the LIBPATH environment variable:

```
export LIBSAVE=$LIBPATH
unset LIBPATH
setupaix.bin -is:javahome /opt/ibm-Java-whatever/java/bin
export LIBPATH=$LIBSAVE
unset LIBSAVE
```

- System i system installation:

Note: The installation of DS CLI on a System i system is done remotely from a Windows, AIX or some other system platform. You cannot install the DS CLI directly on a System i system.

The System i system 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
 - Licensed product 5722-AC3 option *base - Crypto Access Provider 128 bit
 - Licensed product 5722-DG1option *base - IBM HTTP Server for iSeries
 - Licensed product 5722-JV1 options 6 - Java Developer Kit 1.4
 - The latest cumulative PTF package that is installed on the i5/OS
- If you are installing onto a System i system, ensure that the machine that you are installing from is network-attached to the iSeries machine.
- During the installation of the DS CLI application onto a System i system, you must provide the following information:
 - The name of the iSeries machine to which you are installing the DS CLI application.
 - The user name and password that is used to access the designated iSeries machine.

The installation process installs the DS CLI in the following default directory:

AIX /opt/ibm/dscli

HPUX /opt/ibm/dscli

Linux /opt/ibm/dscli

Sun Solaris
/opt/ibm/dscli

Windows
C:\Program Files\IBM\dscli

HP Tru64
/opt/ibm/dscli

iSeries
/ibm/dscli

Novell Netware
SYS:\dscli

Perform the following steps 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 by issuing the following command:

```
mkdir /cdrom -p
```

Create a file system for the CD-ROM by issuing the following command:

```
crfs -v cdrfs -p ro -d cd0 -m /cdrom
```

where *cd0* represents the CD-ROM drive.

Mount the CD-ROM file system by issuing the following command:

```
mount /cdrom
```

HPUX Mount the CD-ROM file system using the path name for your environment by issuing the following commands:

```
ioscan -funC disk | more  
mount /dev/dsk/c?t?d? /<cdrom>
```

Linux Issue the following command on Red Hat systems:

```
mount /dev/cdrom
```

Sun Solaris

Issue 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 Sun 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

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 application on your system using the graphical installation mode.

Users of Windows, Novell, UNIX, and System i systems can install the DS CLI using the graphical mode. Consider the following before you perform the installation of the DS CLI application:

- You cannot install the DS CLI on a computer that contains the ESS CLI.
- You cannot install the DS CLI on a Windows 64-bit operating system.
- You must have a version of Java 1.4.1 or higher 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 1.4.1 or higher exists on your system. If this version is not found in the standard directories, the installation fails.
The DS CLI installation CD contains the Java 1.4.2 version. You can install this version from the CD and continue with the installation of the DS CLI application.
- If the DS CLI application was installed on your client or host machine in the past, ensure that you end any active sessions of the DS CLI before you run the DS CLI installation CD.
- If you are installing onto a Novell system, you are directed to provide the following information:
 - The directory where your Windows drive is mapped
 - The JAVA HOME directory where the version of Java that you want to use is located.
- System i and i5/OS installation considerations:
 - The latest Java group PTF
 - i5/OS 5722-SS1 option 34 - Digital certificate manager
 - Licensed product 5722-AC3 option *base - Crypto Access Provider 128 bit
 - Licensed product 5722-DG1option *base - IBM HTTP Server for iSeries
 - Licensed product 5722-JV1 options 6 - Java Developer Kit 1.4
 - The latest cumulative tape (CUM) package that is installed on the i5/OS
 - The machine that you are installing from must be connected to the i5/OS through an IP connection.

You can install the DS CLI application 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. 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 DS CLI, your old DS CLI sessions might be unusable.

Perform the following steps to install the DS CLI application using the graphical mode:

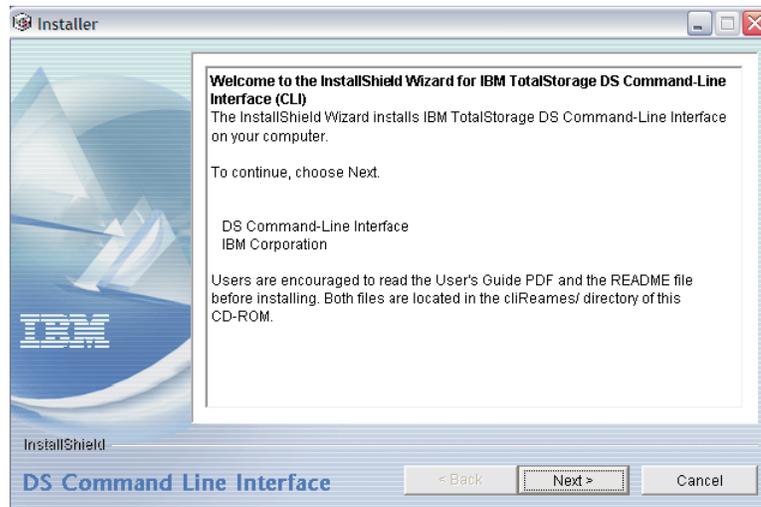
1. Start the setup file that is appropriate for your operating system.
If you are ultimately installing the DS CLI onto a System i system, add **-os400** to your command that includes the setup file. For example, `setupwin32.exe -os400`
Initially the DS CLI installer checks your standard directories for the correct version of Java. If the correct version of Java is not found, you receive the following message:
A suitable JVM could not be found. Please run the program again using the option `-is:javahome <JAVA HOME DIR>`

If you receive this message, see “Correcting the JVM-not-found error” on page 46.

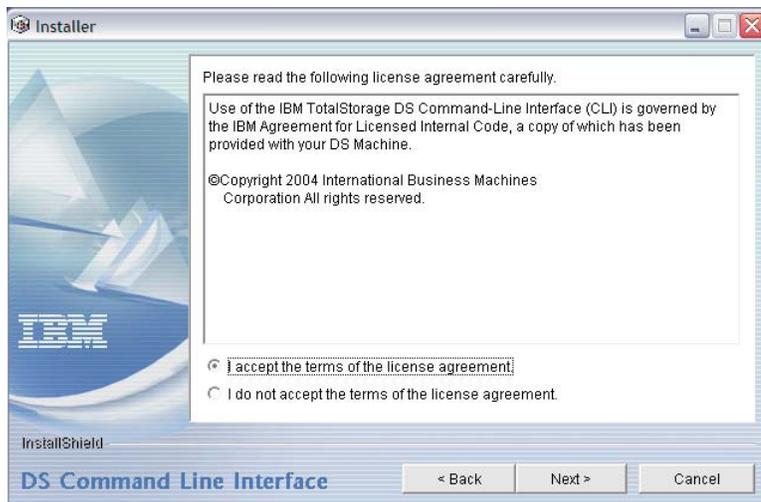
Go to Step 2 if you are installing for a System i system.

Go to Step 3 for all installations other than System i.

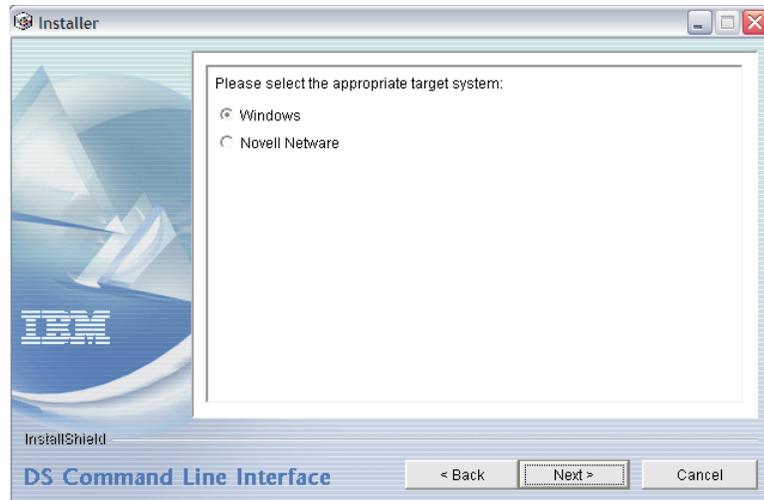
2. The installation wizard displays a sign-on panel that directs you to provide the name of the System i machine to which you are installing the DS CLI application. You must provide the IP address or the DNS name of the i5/OS server. In addition, you must also supply the user name and password that is used to access the designated System i machine. Click **OK**.
3. The Welcome window opens. Click **Next** to continue or **Cancel** to exit the installation.



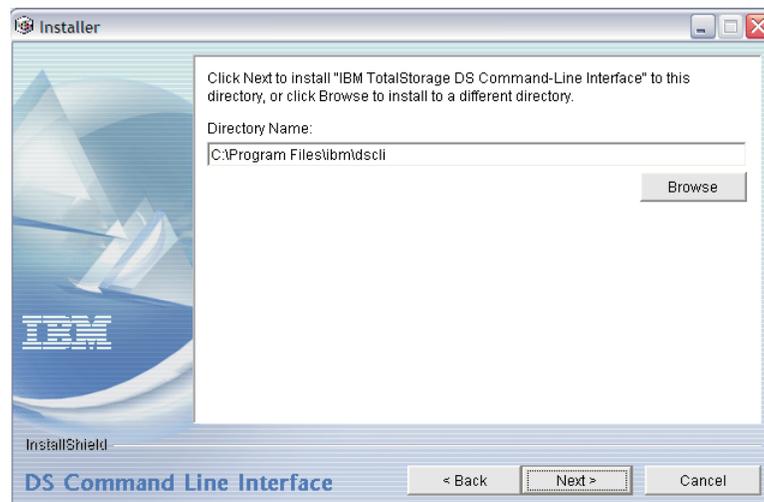
4. The License Agreement window opens. Select “**I accept the terms of this license agreement**” to continue. Click “**I do not accept the terms of this license agreement**” or **Cancel** to exit the installation.



5. The Target System window opens. Select the appropriate target system where you want the DS CLI installed, and then click **Next** to continue or **Cancel** to exit the installation.

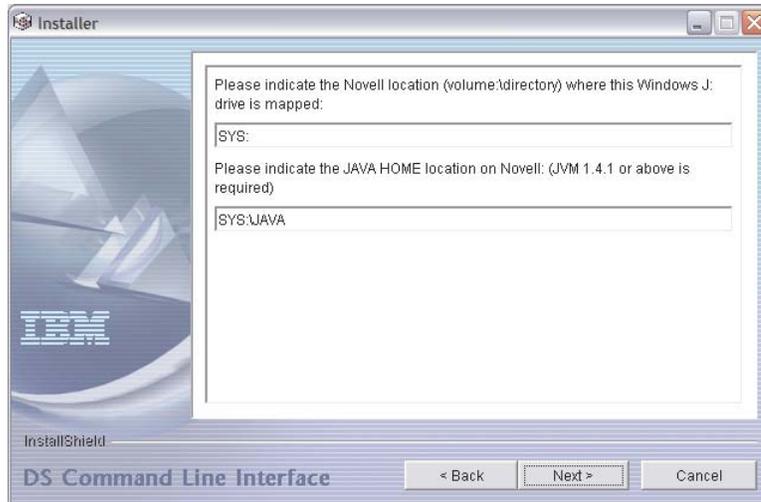


6. The Installation Location window (Windows has been selected as your operating system) opens. If you selected Novell, go to Step 7. If you are installing on a Linux system, go to Step 9 on page 34. If you are satisfied with the Windows default directory, click **Next**. If not, type the directory path where you want to install the DS CLI and click **Next**. Do not use blank spaces in the installation path. Click **Cancel** if you want to exit the installation.

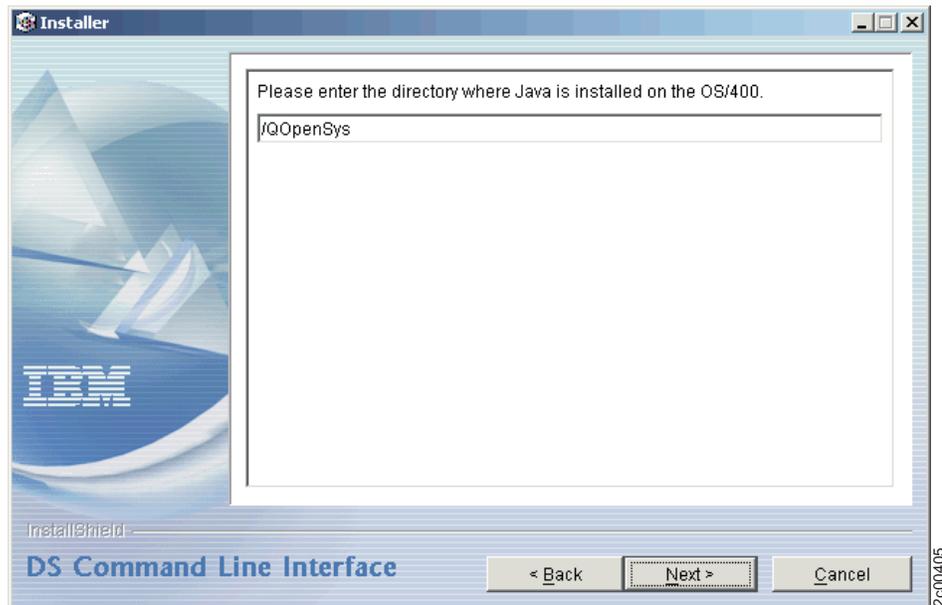


Note: If you are installing onto a System i system, a window that asks for the directory where Java is installed on the i5/OS is displayed when you click **Next**. Go to Step 8 on page 34.

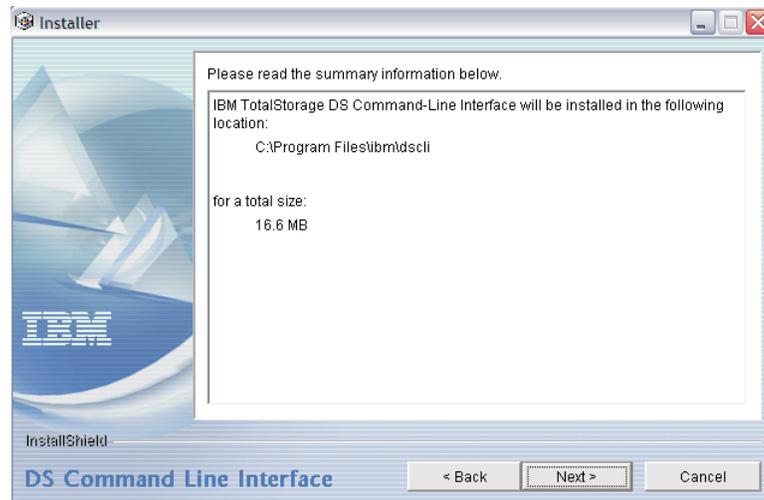
7. The Novell Location window opens. Complete the information for where the Windows drive is installed and where the JAVA HOME directory that contains the version of Java you want to use is located. Click **Cancel** if you want to exit the installation.



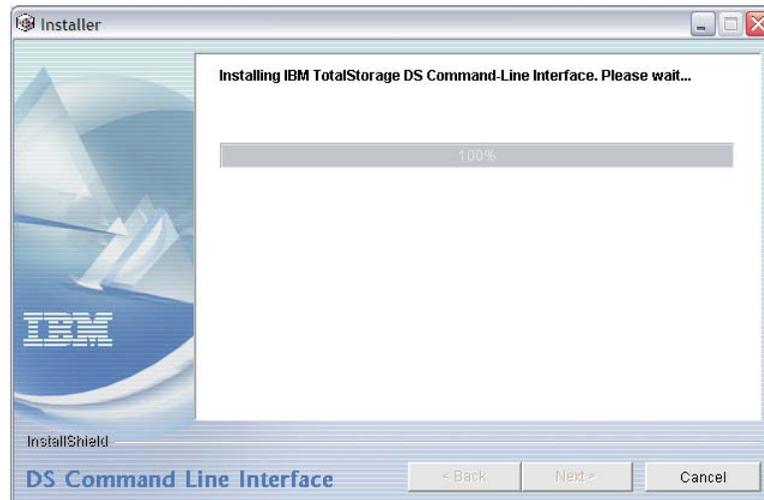
8. The Java directory window for the System i system opens. Insert the directory information. Click **Next** to continue or **Cancel** to exit the installation.



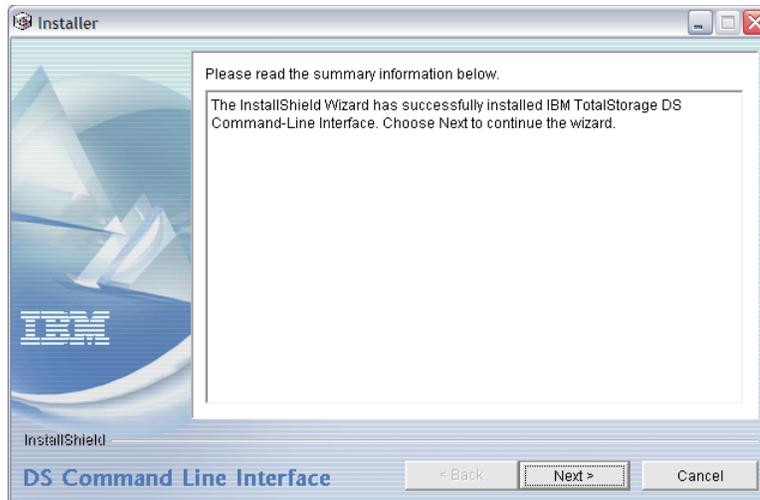
9. The Installation Preview window opens. This window provides the location where the command-line interface will be installed and specifies how much space it will occupy on your drive. Click **Next** to continue or **Cancel** to exit the installation. You can change the installation directory by clicking **Back**.



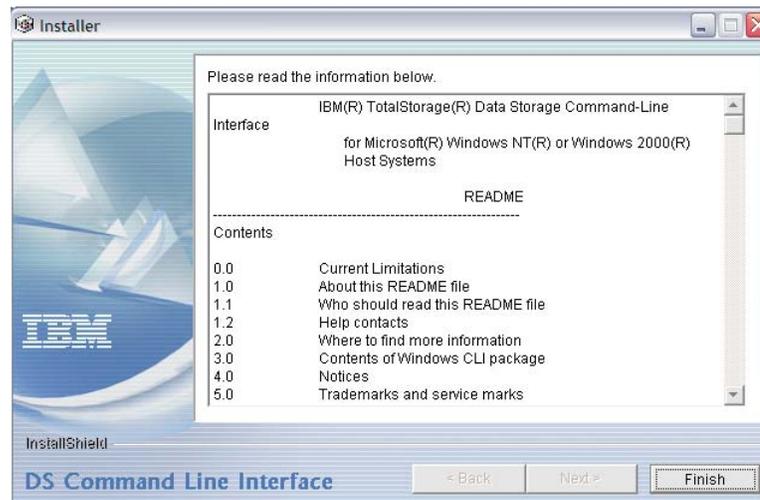
10. The Installation Progress window opens. This window provides the progress of the command-line interface installation. Click **Next** to continue or **Cancel** to exit the installation.



11. The Installation Summary window opens. This window provides the installation summary information. Click **Next** to continue or **Cancel** to exit the installation.



12. The Installation Complete window opens. A reminder is provided in text to read the Readme file. Click **Finish** to complete the installation.



Notes:

1. You can verify that the command-line interface has installed correctly by reviewing the CLI.CFG file in the lib/ directory.
2. The DS CLI application is installed in the following two places in i5/OS:
 - IFS directory *IBM/DS_CLI*. This directory contains the profiles, exe files, java jars, readme files, and so forth.
 - The QDSCLI library. This library contains executable code.
3. Before you can invoke the DS CLI application from the i5/OS, you must add the QDSCLI library to the i5/OS library list.

Installing the DS CLI using the console mode

Complete this task to install the DS CLI application 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. Users can run the installer from a command prompt on a Windows operating system.

Consider the following before you perform the installation of the DS CLI application:

- You cannot install the DS CLI on a computer that contains the ESS CLI.
- You cannot install the DS CLI on a Windows 64-bit operating system.
- You must have a version of Java 1.4.1 or higher 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 1.4.1 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 JVM-not-found error” on page 46*.

The DS CLI installation CD contains the Java 1.4.2 version. You can install this version from the CD and continue with the installation of the DS CLI application.

- If the DS CLI application was installed on your client or host machine in the past, ensure that you end any active sessions of the DS CLI before you run the DS CLI installation CD.
- If you are installing onto a Novell 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
- If you are installing the DS CLI onto an OS/400 system, ensure that the machine that you are installing from is network-attached to the OS/400 machine.
- During the installation of the DS CLI application onto an OS/400 system, you must provide the following information:
 - The name of the OS/400 machine to which you are installing the DS CLI application.
 - The user name and password that is used to access the designated OS/400 machine.

The console mode is primarily used for installations on a Linux operating system or on a UNIX operating system that does not have an X display. Users can run the installer from a command prompt on a Windows operating system. The console mode installer displays its various screens as text.

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.

Perform the following steps to install the DS CLI using the console mode:

1. Insert the DS CLI installation CD into the CD drive. If the Installshield starts, click the **Exit** selection to exit.
2. Open a command prompt and locate the `setupwin32console.exe` or the `setuplinux.bin` file on the DS CLI CD.
3. Type the following command on the command line (unless you are installing onto an OS/400 system): `setup platform<.exe | .bin | .sh> -console`. The Welcome screen is displayed.

For example, for Windows, type: `setupwin32console.exe` or, for Linux, type: `setuplinux.bin -console`

For an installation onto an OS/400 system from a Windows operating system, type: `setupwin32console.exe -os400`

Notes:

- a. When you issue the setupwin32console.exe command on a Windows operating system, you do not need to include the **-console** parameter as part of your command.
- b. 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:
A suitable JVM could not be found. Please run the program again using the option `-is:javahome <JAVA HOME DIR>`
If you receive this message, see “Correcting the JVM-not-found error” on page 46.

```

Initializing InstallShield Wizard...
Searching for Java(tm) Virtual Machine...
.
Searching for Java 1.4.1 by IBM Corporation
Verifying Java 1.4.1 by IBM Corporation
.....
-----
Welcome to the InstallShield Wizard for IBM System Storage DS Command-Line
Interface (CLI)
The InstallShield Wizard installs IBM System Storage DS Command-Line Interface
on your computer.

To continue, choose Next.

DS Command-Line Interface
IBM Corporation

Users are encouraged to read the User's Guide PDF and the README file before
installing. Both files are located in the cliReadmes/ directory of this
CD-ROM.

Press 1 for Next, 3 to Cancel or 4 to Redisplay [1]

```

4. Press 1 for Next. The License agreement screen is displayed. You can also press 3 to Cancel, or 4 to Redisplay.

Please read the following license agreement carefully.

Use of the IBM System Storage DS Command-Line Interface (CLI) is governed by the IBM Agreement for Licensed Internal Code, a copy of which has been provided with your DS Machine.

(C) Copyright 2004, 2005 International Business Machines Corporation All rights reserved.

Please choose from the following options:

[] 1 - I accept the terms of the license agreement.
[] 2 - I do not accept the terms of the license agreement.

To select an item enter its number, or 0 when you are finished: [0] 1

[X] 1 - I accept the terms of the license agreement.
[] 2 - I do not accept the terms of the license agreement.

To select an item enter its number, or 0 when you are finished: [0] 1

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

5. Press 1 to accept the terms of the license agreement. A confirmation prompt is returned with an "X" next to the acceptance line and a 0 in the finished line.
 - a. Press Enter to notify the installer that you are finished with this section. The next action line of instructions is displayed.
 - b. Press 1 for Next. One of the following results occurs depending on the system that you are using to install the DS CLI:
 - The Windows User's screen is displayed if you are installing on a Windows operating system or are pointing to a Novell system. Continue with the next step.
 - The Install Location screen is displayed if you are installing on a Linux operating system. Continue with Step 10 on page 41.You can also press 3 to Cancel, or 4 to Redisplay.

Please select the appropriate target system:

[X] 1 - Windows
[] 2 - Novell Netware

To select an item enter its number, or 0 when you are finished: [0]

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

6. If you want the Windows operating system, press 1. If you want the Novell operating system, press 2. If you choose 1 or 2, a confirmation prompt is returned with an "X" next to the system that you choose and with a 0 in the finished line.
 - a. Press Enter to notify the installer that you are finished with this section.
 - b. Press 1 to continue. The Install Location screen is displayed. This screen can be different depending whether you choose Windows, Novell, or are installing for an OS/400.
 - Continue to Step 7 on page 40 if you choose Windows.
 - Continue to Step 8 on page 40 if you choose Novell.

- Continue to Step 9 on page 41 if you are installing for an OS/400.

You can also press 2 to go to the previous screen, 3 to Cancel, or 4 to Redisplay.

7. (Windows and Linux only) Observe the default directory path that is displayed in the Install Location screen and take the following action (example: the Install Location screen at the end of this step):
 - a. Press 1 to continue if you are satisfied with the default directory.
 - b. Change the directory path where you want to install the CLI (do not use blank spaces in the path) if you are not satisfied, and press 1 to continue.
 - c. Go to Step 10 on page 41 after pressing 1.

The Install Location screen is displayed after pressing 1.

You can also press 2 to go to the previous screen, 3 to Cancel, or 4 to Redisplay.

```
IBM System Storage DS Command-Line Interface Install Location
Please specify a directory or press Enter to accept the default directory.
Directory Name: [C:\Program Files\ibm\dsccli]or [opt/ibm/dsccli]
Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]
```

8. (Novell only) Observe the default directory paths that are displayed in the Installation Location screen for Novell and perform the following action (example: the Install Location screen at the end of this step):

Notes:

- a. SYS represents where the Windows drive is mapped in a Novell system. However, this could also be SYS:\DS\CLI or a volume other than SYS:
 - b. J: (could be any drive letter)
- a. Press 1 to continue if you are satisfied with the default values.
 - b. Enter the applicable value if you are not satisfied with the default, and press 1 to continue.
 - c. Go to Step 10 on page 41 after you press 1.

The Install Location screen is displayed after you press 1.

You can also press 2 to go to the previous screen, 3 to Cancel, or 4 to Redisplay.

```
IBM System Storage DS Command-Line Interface Install Location
Please indicate the Novell location (volume:\directory) where this window J:
drive is mapped:
Enter a value: [SYS:]
Please indicate the JAVA HOME location on Novell: (JVM 1.4.1 or above is
required)
Enter a value: [SYS:\JAVA]
Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]
```

9. (Windows and OS/400 only) Record the default directory path that is displayed in the Install Location screen for OS/400 and perform the following action (example: the Install Location screen shown at the end of this step):
 - a. Press 1 to continue if you are satisfied with the default value.
 - b. Enter the applicable value if you are not satisfied with the default, and press 1 to continue.
 - c. Go to Step 10 after you press 1.

The Install Location screen is displayed after you press 1.

You can also press 2 to go to the previous screen, 3 to Cancel, or 4 to Redisplay.

Please indicate the directory where Java is installed on the OS/400

Enter a value: [/QOpenSys]

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

10. Verify that the installation directory shown on the Install Location screen is the one you want the DS CLI application installed in. Press 1 to continue. The Installation Progress screen is displayed.

Note: Depending on the current configuration of your system, the total size of the installation can vary.

You can also press 2 to go to the previous screen, 3 to Cancel, or 4 to Redisplay.

IBM System Storage DS Command-Line Interface will be installed in the following location:

C:\Program Files\ibm\dscli

for a total size:

16.6 MB

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]

11. Observe the installation progress and wait until the DS CLI installation is completed before pressing 1. Press 1 after the installation completes (there is no reminder to press 1 when the installation is completed). The Install Completed screen is displayed after you press 1.

```

Installing command-line Interface. Please wait...

Installing IBM System Storage DS Command-Line Interface. Please wait...

|-----|-----|-----|-----|
0%      25%     50%     75%     100%
||||||||||||||||||||||||||||||||||||||||

|-----|-----|-----|-----|
0%      25%     50%     75%     100%
||||||||||||||||||||||||||||||||||||||||

Creating uninstaller...

```

12. Press 1 to continue if you are satisfied with the installation. The Readme screen is displayed (not shown).
You can also press 3 to Cancel, or 4 to Redisplay.

```

The InstallShield Wizard has successfully installed IBM System Storage DS
Command-Line Interface. Choose Next to continue the wizard.

Press 1 for Next, 3 to Cancel or 4 to Redisplay [1]

```

13. Press Enter to read the readme information or press 3 to finish the installation process. If you press Enter, you must press q to quit the readme file and then press 3 to finish the installation process. The Readme screen contains specific information about the DS CLI for your operating system.

Note: If you are installing on a Windows operating system, you might have to restart your system. If this is the case, the following screen is displayed and you must complete the next step.

```

To complete the installation you must restart your computer.

[X] 1 - Yes, restart my computer.

[ ] 2 - No, I will restart my computer at a later time.

To select an item enter its number, or 0 when you are finished:

```

14. Select the restart computer item that you want and press 3 to finish the installation process. Use of the DS CLI application becomes effective with a reboot of your system.

You can verify that the command-line interface has installed correctly by reviewing the CLI.CFG file in the lib subdirectory of the install directory that you defined in Step 6.

Installing the DS CLI using the unattended (silent) mode

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

Consider the following before you perform the installation of the DS CLI application:

- You cannot install the DS CLI on a computer that contains the ESS CLI.
- You cannot install the DS CLI on a Windows 64-bit operating system.
- You must have installed a version of Java 1.4.1 or higher on your system in a standard directory. The DS CLI installer checks the standard directories to determine if a version of Java 1.4.1 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 JVM-not-found error” on page 46.*

The DS CLI installation CD-ROM contains the Java 1.4.2 version. You can install this version from the CD and continue with the installation of the DS CLI application.

- If the DS CLI application was installed on your client/host machine in the past, ensure that you end any active sessions of the DS CLI before you run the DS CLI installation CD.
- If you are installing onto a Novell 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.
- If you are installing onto a System i system, ensure that the machine that you are installing from is network-attached to the iSeries machine.
- During the installation of the DS CLI application onto an System i system, you must provide the following information:
 - The name of the iSeries machine to which you are installing the DS CLI application.
 - The user name and password that is used to access the designated iSeries machine.

The unattended (silent) mode of installation allows you to perform the installation from the command line without prompts or feedback. You can create an options/response file and use a text editor to change the default selections to a selection of your choosing.

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 DS CLI, your old DS CLI sessions might be unusable.

Perform the following steps to install the DS CLI application 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. If the InstallShield starts, click the **Exit** selection to exit.
3. Type the following command for all systems except the System i system at the command prompt `setup<platform>.<exe | bin | sh> -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:

```
A suitable JVM could not be found. Please run the program again using the option -is:javahome <JAVA HOME DIR>
```

If you receive this message, see “Correcting the JVM-not-found error” on page 46.

4. Generate the template response file if you must modify the defaults. Two example response files are provided in the cliReadmes directory of the CD: one for Windows, Novell, and iSeries installations, and one for all UNIX and Linux installations. You can copy the applicable response file to a directory of your choice.

The following output is an example of the Windows, Novell, or iSeries response or options file.

```

InstallShield Options File Template
#
# Wizard name: Install
# Wizard source: setup.jar
# Created on: Mon May 09 16:35:04 MST 2005
# Created by: InstallShield Options File Generator
#
# This file can be used to create an options file (i.e., response file) for the
# wizard "Install". Options files are used with "-options" on the command-line
# to modify wizard settings.
#
# The settings that can be specified for the wizard are listed below. To use
# this template, follow these steps:
#
# 1. Enable a setting below by removing leading '###' characters from the
# line (search for '###' to find settings you can change).
#
# 2. Specify a value for a setting by replacing the characters '<value>'.
# Read each setting's documentation for information on how to specify its
# value.
#
# 3. Save the changes to the file.
#
# 4. To use the options file with the wizard, specify -options <file name>
# as a command-line argument to the wizard, where <file name> is the name
# of this options file.
#
#####

#####

#
# License Agreement State
#
# The initial state of the License Agreement panel. Legal values are:
#
# 0 - Nothing will be selected
# 1 - "I accept the terms of the license agreement." will be selected
# 2 - "I do not accept the terms of the license agreement." will be
# selected
#
# For example, to configure the panel to initially display "I accept the
# terms of the license agreement.", use
#
# -W license.selection=2

-W license.selection=1

#####

#
# User Input Field - Win_Novell_select
#
# This field only has meaning if you are running the installer on Windows.
# You must choose whether it should install the Windows CLI or the Novell
# CLI.
#
# If you choose Novell, then make sure that the installLocation (see below)
# is on a Novell drive.
#
# 1 - Novell Netware
# 2 - Windows
#
-W win_user_input.Win_Novell_select="2"

#####

```

```

IBM System Storage DS command-line Interface Install Location
#
# The install location of the product. Specify a valid directory into which the
# product should be installed. If the directory contains spaces, enclose it in
# double-quotes. For example, to install the product to C:\Program Files\My
# Product, use
#
### -P installLocation="C:\Program Files\My Product"
#
-P installLocation="C:\Program Files\IBM\dscli"

#####
#
# User Input Field - os400_java_location
#
#
# This is only used on i5/OS iSeries system installations.
#
#
### -W os400_java_location_panel.os400_java_location="<value>"
#
#####
#
# User Input Field - novell_location_mapped
#
# This is only used on Novell installations.
#
#
### -W novell_user_input.novell_location_mapped="<value>"
#
#####
#
# User Input Field - novel_java_location
#
# This is only used on Novell installations.
#
#
### -W novell_user_input.novel_java_location="<value>"
#
#####

```

5. Save the modifications to the response file.
6. Type the following command for all systems (including UNIX and Linux) except the i5/OS iSeries system at the command prompt `setup<platform>`. `<exe | bin | sh> -silent -options options.txt` and press the **Enter** key on your keyboard to start the installation process in unattended (silent) mode.
7. Type the following command if you are installing the DS CLI onto an i5/OS iSeries system `setup<platform>`. `<exe | bin | sh> -os400 -silent -options options.txt` and press the **Enter** key on your keyboard to start the installation process in unattended (silent) mode.

You can verify that the command-line interface has installed correctly by reviewing the CLI.CFG file in the lib/ subdirectory of the installation directory.

Correcting the JVM-not-found error

Complete this task to correct the JVM-not-found error.

The JVM-not-found error occurs if the DS CLI installer cannot find the correct version of Java in the standard directories of your system. You must have Java 1.4.1 or higher on your system for the DS CLI application to work.

Notes:

1. This error might also occur if you are installing the DS CLI on an AIX system. The LIBPATH environment variable can interfere with the

installation of the DS CLI and results in the display of the JVM-not-found error. It is recommended that you 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.

2. If you have already installed the DS StorageManager, you have installed Java 1.4.2 on your system. Because it was not installed in one of your standard directories, the DS CLI installer did not find it during its initial check.

If Java 1.4.1 is not found during the initial check, the following message is displayed:

```
A suitable JVM could not be found. Please
run the program again using the option -is:javahome <JAVA HOME DIR>
```

Perform either step 1 or step 2 of the following steps to correct the JVM not found error (see Step 4 for AIX):

1. (Optional) Run the installer again, and provide the path to the JVM. The following examples represent directories where you can find the correct version of Java:

For a Windows system, specify the following:

```
setupwin32 -is:javahome
C:\Program Files\IBM\dsniserver\java\jre
```

For a Linux system, specify the following:

```
setuplinux.bin -is:javahome
/opt/ibm-Java-whatever/java/bin
```

Note: If you use this argument, the installer attempts to use whatever JVM that you specify, even if it is an unsupported version.

Go directly to step 4 after this step completes.

2. (Optional) Open the DS CLI installation CD and install the Java 1.4.2 version into one of your standard directories.
3. Continue with the installation of the DS CLI application.
4. (AIX) Run the following commands to sequentially disable the LIBPATH environment variable, install the DS CLI, and restore the LIBPATH environment variable:

```
export LIBSAVE=$LIBPATH
unset LIBPATH
setupaix.bin -is:javahome /opt/ibm-Java-whatever/java/bin
export LIBPATH=$LIBSAVE
unset LIBSAVE
```

Proceed to the next step after you disable the LIBPATH environment variable.

Installing DS CLI on an OpenVMS system

The DS Command-Line Interface (CLI) provides a full function command set that allows you to check your storage unit configuration and perform specific application functions when necessary. It also enables open systems hosts to invoke and manage FlashCopy, Metro and Global Mirror functions through batch processes and scripts.

This topic describes how to install the DS CLI on an OpenVMS system. Additionally, it provides hints for integrating the DS CLI into the operating system environment.

Preparing to install the DS CLI on an OpenVMS System

The IBM System Storage DS CLI program enables open systems hosts to invoke and manage FlashCopy and Metro and Global Mirror functions through batch processes and scripts. The DS CLI is installed on an OpenVMS system by using the Polycenter Software Installation utility.

Be aware of the following requirements before you begin the installation of the DS CLI:

- The installation process places all command-line interface files in the [IBMDCLI...] directory tree as a subdirectory of the Polycenter destination directory. You can specify this directory by using the **PRODUCT INSTALL** command with the **/DESTINATION=devicename:[directoryname]** qualifier. If you omit the device name, the Polycenter installation utility uses your current default device. If you omit the directory name, the Polycenter Software Installation utility uses the [VMS\$COMMON] directory as the default destination directory.

If you do not use the **/DESTINATION** qualifier at all, the utility installs the software in the location that is defined by logical name **PCSI\$DESTINATION**. If this logical name is not defined, the utility installs the software in **SYSSYSDEVICE:[VMS\$COMMON]**.

- Extract and check the Release Notes file from the DS CLI installation package in the root directory of the installation CD.
- Ensure that you have at least 140000 blocks (approximately 70 MB) of free space on the installation disk.
- Install the required prerequisite patches operating system patches (ECOs) on your OpenVMS host system.

The command-line interface installation process automatically installs Java™ Runtime Environment (JRE) 1.4.2-4 on your host system. The JRE requires several ECOs. For detailed patch information, see the Java SDK v1.4.2 patch installation page at:

http://h18012.www1.hp.com/java/download/ovms/1.4.2/sdk1.4.2_patches.html

To download these ECOs, use the HP IT Resource Center (ITRC) database to perform a search for the patches from which you can select. These download steps are described at the Java SDK v1.4.2 patch installation page.

Installing the DS CLI on your OpenVMS system

Use the following steps to install the DS CLI application in your OpenVMS environment.

Before you install the DS CLI application, verify that the prerequisite software and hardware are installed on your system. The installation program checks for prerequisites and stops if any prerequisites are missing.

1. Log on to your host system as a user with SYSLCK, SYSNAM, SYSPRV, (or a system UIC), TMPMBX, and CMKRNL privileges.
2. Insert the DS CLI product CD-ROM into the CD drive.

- Mount the CD drive. For example, for an IDE CD device DQA0, type the following command: **MOUNT /NOASSIST /OVERRIDE=IDENTIFICATION /MEDIA_FORMAT=CDROM DQA0**: A message similar to the following is displayed.

```
%MOUNT-I-WRITELOCK,volume is write locked
%MOUNT-I-CDROM_ISO, : (1 of 1) , mounted on VMS1$DQA0:
```

For a SCSI CD device DKA *nnn*, type the following command: **MOUNT /NOASSIST /OVERRIDE=IDENTIFICATION /MEDIA_FORMAT=CDROM DKA_{nnn}**: where *nnn* represents the number that is assigned by the OpenVMS system to your CD device.

- Type the following command: **DIRECTORY/FULL DQA0:[000000]IBM-AXPVMS-DSCLI-* .PCSI** and press **Enter** to access the command-line interface installation package in the root directory of the CD. Output similar to the following is displayed.

```
Directory DQA0:[000000]
IBM-AXPVMS-DSCLI-V0500-01F96-1.PCSI;1   File ID: (4,7,0)
Size:          55.79MB/55.80MB   Owner:   [0,0]
Created:       9-MAR-2005 04:07:22.25
Revised:       9-MAR-2005 04:09:43.98 (1)
Expires:       None specified
Backup:        No Backup record
Effective:     None specified
Recording:     None specified
Accessed:     None specified
Attributes:    None specified
Modified:     None specified
Linkcount:    1
File organization: Sequential
Shelved state: Online
Caching attribute: Writethrough
File attributes: Allocation: 114282, Extend: 0, Global buffer count: 0
                  Version limit: 0, Backups disabled
Record format: Undefined, maximum 0 bytes, longest 0 bytes
Record attributes: None
RMS attributes: None
Journaling enabled: None
File protection: System:RWED, Owner:RWED, Group:RWED, World:RWED
Access Cntrl List: None
Client attributes: None

Total of 1 file, 55.79MB/55.80MB
```

- Type the following command and press **Enter** to extract the command-line interface for OpenVMS release notes: **PRODUCT EXTRACT RELEASE_NOTES DSCLI /SOURCE=DQA:[000000] /FILE=filespec**

Note: If you do not use the **/FILE** qualifier, the release notes are written to the DEFAULT.PCSI\$RELEASE_NOTES file in your current default directory. Read the release notes before continuing to ensure that you are equipped with the information that you need to complete your installation successfully.

- Type the following command and press **Enter** to invoke the command-line interface installation process: **PRODUCT INSTALL DSCLI /SOURCE=DQA0:[000000] /DESTINATION=devicename:[directoryname] /RECOVERY_MODE**. A message similar to the following is displayed.

The following product has been selected:

IBM AXPVMS DSCLI V5.0-1F96 Layered Product

Do you want to continue? [YES]

The **/DESTINATION** qualifier can be omitted from the command. However, IBM recommends that you use the optional **/RECOVERY_MODE** qualifier. For a detailed description of all **PRODUCT INSTALL** command qualifiers and parameters, see the HP OpenVMS System Management Utilities Reference Manual or the OpenVMS online help.

7. Press Enter to continue the installation process. The following message is displayed.

Configuration phase starting ...

You will be asked to choose options, if any, for each selected product and for any products that may be installed to satisfy software dependency requirements.

IBM AXPVMS DSCLI V5.0-1F96: IBM DS Command Line Interface (DS CLI)

(C) Copyright IBM Corp. 1999, 2005. All Rights Reserved.

International Business Machines Corporation (IBM)

No PAK

Do you want the defaults for all options? []

8. Type No and press Enter to review the installation options.

Note: You can type Yes if you have decided to install with the default options.

The following CLI license message is displayed after typing Yes or No.

License Agreement IBM System Storage DS Command-line Interface

Copyright 2005 International Business Machines Corporation
All rights reserved.

Use of the IBM System Storage DS Command-line Interface (DS CLI) is governed by the IBM Agreement for Licensed Internal Code, a copy of which has been provided with your DS Machine.

Answer "Yes" to install the DS CLI code. By installing this code, you are certifying that you have read and accept the IBM DS CLI License agreement.

Answer "No" to terminate and exit the IBM DS CLI installation.

Do you want to continue? [YES]

9. Press Enter to accept the DS CLI license agreement. The following Java license message is displayed.

License Agreement HP Java Run-Time Environment for OpenVMS

The IBM DS CLI requires the Java 1.4.2 Java Runtime Environment (JRE). Installing the DS CLI program automatically installs the required JRE. The installed JRE is the intellectual property of and licensed by Hewlett-Packard Company.

You can view the license agreement on the World Wide Web at:

http://h18012.www1.hp.com/java/download/ovms/1.4.2/rte1.4.2_license.html

Answer "Yes" to install the Java code. By installing this code, you are certifying that you have read and accepted the HP Java License agreement.

Answer "No" to terminate and exit the IBM DS CLI installation.

Do you want to continue? [YES]

10. Press Enter to accept the Java license agreement. The command-line interface requires Java 1.4.2 on an OpenVMS Alpha host system. If you answer No, the installation process automatically ends and exits. The following library update warning message is displayed.

WARNING: By default the system Help and Message libraries will be updated.

The IBM DS CLI program provides local Help and Help Message library files. By default, the CLI installation integrates these local libraries into the OpenVMS system Help and Help Message libraries.

To prevent the system libraries from being modified chose to review the installation options and answer "No" when prompted to update the libraries.

Do you want to continue? [YES]

11. Press Enter to continue the installation process. The following configuration question is displayed.

Would you like the local IBM DS CLI Help and Help Message libraries to be integrated into the OpenVMS system libraries?

If you answer "Yes", the following OpenVMS libraries will be updated:

SYS\$COMMON:[SYSHLP]HELPLIB.HLB
SYS\$COMMON:[SYSHLP]MSGHLP\$LIBRARY.MSGHLP\$DATA

If you answer "No", OpenVMS system libraries will not be modified.

In every case, local libraries are available under:

IBMDSCLI\$HELP:IBMDSCLI_OVR.HLB
IBMDSCLI\$HELP:IBMDSCLI_MESSAGES.MSGHLP\$DATA

Modify the system Help and Help Message libraries? [YES]

12. Press Enter to confirm the library update option (or type No and press Enter to deny the library update option). The following confirmation message is displayed:

Do you want to review the options? [N0]

13. Press Enter to confirm and accept all selections. The following installation message with completion status is displayed.

```
Execution phase starting ...

The following product will be installed to destination:
  IBM AXPVMS DSCLI V5.0-1F96
  DISK$V732_ALPHA:[VMS$COMMON.]

Portion done: 0%...10%...20%...30%...60%...70%...80%...90%...100%

The following product has been installed:
  IBM AXPVMS DSCLI V5.0-1F96          Layered Product
```

14. Review the Installation Verification Procedure (IVP) report, which is similar to the following output, and check for possible errors.

```
%PCSI-I-IVPEXECUTE, executing test procedure for IBM AXPVMS DSCLI
V5.0-1F96 ...
dscli -ver
IBM DSCLI Version: 5.0.1.96
%PCSI-I-IVPSUCCESS, test procedure completed successfully
```

15. Ensure that the installation completes. When the Polycenter Software Installation utility finishes the command-line interface installation process, a message similar to the following is displayed.

```
IBM AXPVMS DSCLI V5.0-1F96: IBM DS Command Line Interface (DS CLI)

Insert the following lines in SYS$MANAGER:SYSTARTUP_VMS.COM:
  @PCSI$DESTINATION:[IBMDSCLI.MGR]IBMDSCLI$STARTUP.COM
Insert the following lines in SYS$MANAGER:SYSHUTDOWN.COM:
  @IBMDSCLI$MANAGER:IBMDSCLI$SHUTDOWN.COM

Users of this product require the following lines in their login
command procedure:
  @IBMDSCLI$MANAGER:IBMDSCLI$LOGIN.COM

Release notes for IBM DS CLI available in IBMDSCLI$HELP
```

16. Unmount the CD drive and remove the CD-ROM.

The command-line interface provides program startup, login, and shutdown procedures in the [destinationdir.IBMDSCLI.MGR] directory. The installation process runs the startup and login procedures immediately before invoking the IVP procedure. But for persistent setup, you must integrate the startup, login, and shutdown procedures.

Using the DS CLI on an OpenVMS system

To use the DS CLI on an OpenVMS system, you must integrate the required `IBMDSCLI$STARTUP.COM` and `IBMDSCLI$LOGIN.COM` procedures. You can also integrate the optional `IBMDSCLI$SHUTDOWN.COM` procedure. This integration is accomplished when you use the OpenVMS persistent setup procedure.

You must add the required `IBMDSCLI$STARTUP.COM` procedure to your system startup processes. The `IBMDSCLI$STARTUP.COM` procedure defines the logical names that are required for the command-line interface in your system logical name table and installs some images with enhanced privileges. This procedure is intended to be invoked during the system startup.

You also must add the required `IBMDSCLI$LOGIN.COM` procedure to your system login processes. The `IBMDSCLI$LOGIN.COM` procedure sets up the JRE that is required by the command-line interface and defines the DSCLI command as a foreign DCL command symbol. It is intended to be invoked during the system-wide `SYS$MANAGER:SYLOGIN.COM` or user-specific `SYS$LOGIN:LOGIN.COM` procedure.

In order to use the DS CLI interface from start to finish in your OpenVMS system, you can add the optional `IBMDSCLI$SHUTDOWN.COM` procedure to your system shutdown processes. The `IBMDSCLI$SHUTDOWN.COM` procedure performs the removal operations for privileged images and undefines system-wide logical names that are associated with the DS CLI. It is intended to be invoked during the system shutdown process.

1. Integrate the `IBMDSCLI$STARTUP.COM` procedure into your system startup by adding the following line to the `SYS$MANAGER:SYSTARTUP_VMS.COM` script: `$ @destinationdev:[destinationdir.IBMDSCLI.MGR] IBMDSCLI$STARTUP`

destinationdev

Specifies the name of the device that contains the command-line interface installation directory

destinationdir

Specifies the name of the directory where you just installed the command-line interface.

Note: You can alternatively add the `IBMDSCLI$STARTUP.COM` procedure to the `SYSMAN` startup database.

2. Integrate the `IBMDSCLI$LOGIN.COM` procedure into the system-wide or user-specific login by adding the following line:
`$(IBMDSCLI$MANAGER:IBMDSCLI$LOGIN.COM`

Note: Run the `IBMDSCLI$LOGIN.COM` procedure only after you have successfully run the `IBMDSCLI$STARTUP.COM` procedure.

3. Integrate the `IBMDSCLI$SHUTDOWN.COM` procedure by adding the following line to the `SYS$MANAGER:SYSHUTDOWN.COM` script:
`$(IBMDSCLI$MANAGER:IBMDSCLI$SHUTDOWN`

Note: This step is optional. However, processing this step allows your system to make the full use of the DS CLI application.

Removing the DS CLI from an OpenVMS system

The removal of the DS CLI application from an OpenVMS system requires not only a removal from your main system but also a removal from your startup, login, and shutdown processes.

Perform the following steps to remove the CLI from your OpenVMS system:

1. Log on to your host system as a user with SYSLCK, SYSNAM, SYSPRV (or a system group UIC), TMPMBX, and CMKRNL privileges.
2. Type the following command at the command prompt to start the uninstallation process: `PRODUCT REMOVE IBMDSCLI`. A message similar to the following is displayed.

```
The following product has been selected:
  IBM AXPVMS DSCLI V5.0-1F96           Layered Product
Do you want to continue? [YES]
```

3. Press Enter to confirm the uninstallation. The following uninstallation confirmation message with completion status is displayed.

```
The following product will be removed from destination:
  IBM AXPVMS DSCLI V5.0-1F96           DISK$V732_ALPHA:[VMS$COMMON.]
Portion done: 0%...10%...20%...30%...40%...50%...60%...70%...80%...90%...100%
```

When the uninstallation process ends, a message similar to the following is displayed.

```
The following product has been removed:
  IBM AXPVMS DSCLI V5.0-1F96           Layered Product
```

4. Remove the command-line interface startup, login, and shutdown functions from your system startup, login, and shutdown processes.

Completing DS CLI postinstallation

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

You must install the DS CLI before you complete these postinstallation tasks.

After you install the DS CLI application, how do you log into the application? After the initial login, what are some of the first tasks that you must accomplish before you can get the full benefit of the DS CLI application? The instructions in this section answer these questions, as well as, describe how to configure your volumes .

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

1. Initiate the DS CLI to begin using it in either single-shot, script, or interactive command mode.
2. Set up your required user accounts and passwords.
3. Set your DS CLI default configuration settings.
4. Activate your licensed functions. This includes obtaining your feature activation codes and applying the feature activation codes to your storage unit.
5. Configure new fixed block 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.

Using the DS CLI application

Your DS8000 model is ready for you to begin using the DS CLI after it is installed. There are three command modes available to you.

You must log into DS CLI to use the command modes. There are three command modes for the DS CLI:

- Single-shot
- Interactive
- Script

Note: You must not start more than 100 DS CLI sessions simultaneously. Starting more than 100 DS CLI sessions simultaneously can result in connection problems.

Logging into the DS CLI application

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

You must ensure that you are in the directory where you installed the DS CLI application. The following list provides a reminder of the supported operating systems default directories where the DS CLI is installed if the directory designation is not changed:

AIX /opt/ibm/dscli

HPUX /opt/ibm/dscli

Linux /opt/ibm/dscli

Sun Solaris
/opt/ibm/dscli

Windows
C:\Program Files\IBM\dscli

HP Tru64
/opt/ibm/dscli

iSeries
/ibm/dscli

Novell Netware
SYS:\dscli

When you log into the DS CLI application (type `dscli` at the command prompt), you must provide the following information:

HMC1
Specify the primary management console.

Note: If you are using 2105 models as part of your network and are going to use the Copy Services functions, you must specify the IP address of the domain control server where you have installed the DS CLI application.

User Name

Specify the name of the user account. The default account for the first login is **admin**.

Password

Specify the user password. The default password for the admin account is **admin**. However, this password is only good for your first login.

Note: Because the password for the admin account expires 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.

The first time that you log in to the DS CLI, you can specify this information using either of the following two methods:

- Ensure you are in the directory where you installed the DS CLI application and type the `dscli` command at the command prompt. Supply all the log in information with the command. For example: `dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t`.

Use this command when you use the single-shot mode for the first time and when the DS CLI application is not active on your system. In addition, when you use the single-shot mode, you must include the command that you want to process. For example, if you want to process the **lssi** command, if you have not activated the DS CLI application, and if you are using the single-shot mode, type: `dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t lssi`.

- When you log into the DS CLI application (from the directory where you installed the DS CLI application) by typing `dscli`, you are prompted to supply the information for HMC1, user name, and password.

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 hitting the Enter key and then entering the rest of the command on a second line. The DS CLI application 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 issue an occasional command but do not want to keep a history of the commands that you have issued.

You must supply the login information and issue the command that you want to process at the same time. Perform the following steps to use the single-shot mode:

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

```
dscli -hmc1 mtc032h.storage.tucson.ibm.com -user admin -passwd topn0t  
lssi -s -fullid -hdr off
```

Here is an example of this same command in i5/OS without the report delimiters:

```
DSCLI SCRIPT(*NONE) HMC1('9.5.17.156') USER(admin) PASSWORD(itso4all)  
DSCL(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 manufacture name (IBM), the machine type (2107), and the serial number.

Notes:

- a. The command example uses the **-fullid** DS CLI command flag. The **-fullid** command flag 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 flag which turns off the header that is generally associated with the report that is generated from the **lssi** command.
 - c. Almost every DS CLI command requires the use of the storage image ID. If you add your target storage ID into your profile file under the **-devid** designation, you are not required to provide the storage image ID when you issue a command. However, if you do provide the **-dev** (*storage_image_ID*) parameter, the value that you type takes priority over the value contained in your profile file.
2. Wait for the command to process. The following type of report is generated that lists the storage image IDs that are associated with the storage complex.

IBM.2107-75FA111
IBM.2107-75FA111
IBM.2107-75FA120

Using the DS CLI script command mode

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

Consider the following when using 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.

Note: It is not the intent of this instruction to tell you how to write a script. An example script is displayed for your use as a guide.

You can issue 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 using the following format:

```
dscli> -hmc1 mtc032h.storage.tucson.ibm.com -user admin  
-passwd password -script ~/bin/mkpprcpairs
```

Note: If you are using i5/OS and have already logged on to the DS CLI application, you invoke the script mode using the following format:
DSCLI SCRIPT('/myscript') USER(admin) OUTPUT('/outfile')

2. Wait for the script to process and provide a report regarding the success or failure of the process.

The following is an example of a script that could be used to establish remote mirror and copy relationships for volume pairs.

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

Using the DS CLI interactive command mode

Use the DS CLI interactive command mode when you have multiple transactions to process 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.

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

1. Log on to the DS CLI application at the directory where it is installed.

Note: If you make a mistake and type the wrong user name or password, do not try to correct the error within the current session. Exit the DS CLI session that you are in and log in to a new DS CLI session.

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.

Note: Issue the **setoutput** command if you know that you will be using interactive mode for a lengthy period of time. The **setoutput** command allows you to set or display command output format options. You control how the reports that are generated by the list and the show commands are displayed on your computer. 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 (horizontal table) format

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

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

```
[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 retry the last command. Or, issue `dscli>!3` to retry 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 **lssi** or **lssu** DS CLI commands to list the storage image IDs that are associated with your storage complex. It is recommended that you 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, perform the following steps. The storage image ID consists of the manufacture name (IBM), the machine type (2107), and the serial number.

1. Log into the DS CLI application in interactive command mode (if you are not already logged in).
2. Type the following command format at the `dscli` command prompt to obtain the storage image IDs. Substitute **lssi** or **lssu** for the command name variable in the example.

```
dscli>commandname -s -fullid -hdr off
```

Note: The result of using the parameters and command flags as shown in the example is the same for either command. However, each command can produce different information when you use their associated parameters. See the **lssi** or **lssu** commands for more details.

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.2107-75FA111
IBM.2107-75FA120

DS CLI command help

You can obtain online help for each CLI command. To obtain the help, type the word *help* and the command name at the dscli command prompt. There are additional parameters that you can use with the help command that designate the type of help information that you can receive.

The **help** command contains the following 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.
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.

Note: You cannot use the **-s** and **-l** parameters with the following help command flags: **-h**, **-help**, and **-?**.

Examples

The following examples represent the type of information that is displayed when you use various parameters with the help command. Each of these examples start at the dscli command prompt.

Note: 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 -r20`. This command pauses your page listing after 20 entries and prompts you to **press any key to continue**.

Example 1

```
dscli>help
```

This line of input provides the entire list of DS CLI commands. Only the command names are displayed. No other details are provided.

Example 2

```
dscli>help -s
```

This line of input provides the entire list of DS CLI commands and a short description for each command. The short description explains what each command accomplishes.

Example 3

```
dscli>help -l
```

This line of input provides the entire list of DS CLI commands and the syntax for each command. In addition, you see all the help syntax and formatting parameters, which can make your reading of the help difficult. For example, you can see something similar to the following for each command:

```
lsextpool [ { -help|-h|-? } ] [ { -l (long)|-s (short) } ]  
[-fmt default|xml|delim|stanza] [-p on|off] [-delim char]  
[-hdr on|off] [-bnr on|off] [-r #] [-v on|off] [-fullid]  
[-hmc1 SMC1] [-hmc2 SMC2] [-user username] [-passwd password]  
-dev storage_image_ID [-stgtype fb|ckd] [-rankgrp 0|1]  
[Extent_Pool_ID ... | -]
```

The following line is the actual syntax for the **lsextpool** command: -dev storage_image_ID [-stgtype fb|ckd] [-rankgrp 0|1] [Extent_Pool_ID ... | -]

Example 4

```
dscli>lssi -h  
dscli>lssi -help  
dscli>lssi -?  
dscli>help lssi
```

Any of the previous lines of input generates the entire help page that is associated with the **lssi** command. This is the same information that is found in the *IBM System Storage DS8000 Command-Line Interface User's Guide* or in the Information Center for the **lssi** command.

Example 5

```
dscli>help -s lssi
```

This line of input provides the short description that is associated with the designated command. For instance, you might see the following:

```
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 when using the **lsavailprcport** and **mkpprcpath** commands.

Example 6

```
dscli>help -l lssi
```

This line of input displays the entire reference page found in both the *IBM System Storage DS8000 Command-Line Interface User's Guide* and in the Information Center for the **lssi** command.

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: Only the code number is displayed when you are using the single-shot or script mode not the code description.

Code	Category	Description
0	Success	Specifies that the command is successfully processed
2	Syntax Error	Specifies that there is an error in the way that the command is presented (misaligned or wrong parameters) for processing
3	Connection Error	Specifies that there is a connectivity or protocol error
4	Server Error	Specifies that an error occurs during a function call to the application server
5	Authentication Error	Specifies that an error occurs during the authentication process
6	Application Error	Specifies that an error occurs due to a MetaProvider client application specific process

Perform the following steps to obtain, 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.

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 network.

The following licenses can be activated depending on your purchase:

- 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

There are multiple codes that are associated with these features. To obtain the information that you need to activate these licenses and features in your storage unit, go to the IBM Disk Storage Feature Activation (DSFA) Web site at:

<http://www.ibm.com/storage/dsfa/>

Download your codes onto a diskette in XML format. You can then import the codes from the XML file when you process the DS CLI `applykey` command.

Note: 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 need the 2244 serial number (which you can find on the License Function Authorization document).

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

Perform the following steps to activate your license activation codes.

1. Log into the DS CLI application in interactive command mode (if you are not already logged in).
2. Issue the DS CLI **applykey** command at the dscli command prompt as follows (this example presumes that your XML file is named "keys" and it resides on a diskette in your A: drive): dscli> applykey -file a:\keys.xml
IBM.2107-75FA120
3. Press Enter. When the process has completed, the following message is displayed:
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI
Version 5.0.0.0 DS: IBM.2107-75FA120
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 **lskey** command as follows: lskey IBM.2107-75FA120
5. Press Enter and the following type of report is displayed:
Sun Apr 09 02:23:49 PST 2006 IBM DS CLI

Activation Key	Authorization Level (TB)	Scope
Operating Environment (OEL)	45	All
Remote mirror and copy (RMC)	25.1	All
Metro Mirror (MM)	25	All
Global Mirror (GM)	25	All
Metro/Global Mirror (MGM)	25	All
Remote mirror for z/OS (RMZ)	25.1	CKD
Point in time copy (PTC)	On	All
Parallel access volumes (PAV)	On	CKD
IBM HyperPAV	On	CKD

Setting up user accounts using the DS CLI

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

The admin account is set up automatically at the time of installation. It is accessed using the user name **admin** and the default password **admin**. This password is temporary and expires after its initial use. You must change the password before you can use any of the other functions. There are 7 groups the administrator can assign to a user. A user can be assigned to more than one user group. The groups and the associated functions allowed by the assignment are as follows:

admin

All users that you assign to the administrator user group allows access to all management console server service methods and all storage image resources.

op_volume

The volume operator user group 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 group inherits all authority of the monitor user group.

op_storage

The storage operator user group allows access to physical configuration service methods and resources, including storage complex, storage image, array, rank, and extent pool objects. This user group inherits all the authority of the op_copy_services and monitor user groups, excluding security methods.

op_copy_services

The copy services operator user group allows access to all Copy Services service methods and resources, excluding security methods. In addition, this user group inherits all authority of the monitor user group.

service

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

monitor

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

no access

The no access user group does not allow access to any service methods or storage image resources. By default, this user group is assigned to any user account in the security repository that is not associated with any other user group.

In addition to assigning users to one or more user groups, you also must assign a default password to each user. When you notify users of their group 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 **admin** 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 group or groups) and an initial password. If you assign multiple roles to an account, ensure that you separate the different roles by using a comma for example, op_volume, op_storage. See the mkuser command description for more details.

1. Log into the DS CLI application in interactive command mode.
2. 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 testuser
3. Press Enter and observe the processing result. A successful process returns the following display:

```
Sun Aug 11 02:23:49 PST 2004 IBM DS  CLI
Version 5.0.0.0 DS: IBM.2107-75FA120
User Name testuser successfully created.
```

Creating a default configuration setup with a profile file

You can create default settings for the command-line interface by defining one or more profiles on the system. For example, you can specify the management console

(MC) for the session, the output format for list commands, the number of rows per page in the command-line output, and the banner that is included with the command-line output.

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

You have several options for using 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\IBM\DSCLI\profile\dscli.profile` is the directory path for the Windows platform and `/opt/ibm/dscli/profile/dscli.profile` is the directory path for UNIX and Linux platforms.
- 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 the storage unit operations. Save the profile in the user profile directory. For example:

```
<user_home>\dscl\profile\operation_name1
<user_home>\dscl\profile\operation_name2
```

These profile files can be specified using the DS CLI command parameter `-cfg <profile_name>`. If the `-cfg` profile file is not specified, the user's default profile file is used. If a user's profile file does not exist, the system default profile file is used.

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

Windows operating system

For a Windows operating system, the property value defaults to the environment variable `%USERPROFILE%`. As a result, your personal profile is `C:\Documents and Settings\username\dscli\profiles\dscli.profile`.

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/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`.

Note: The values of the Java system properties can be redefined by JRE options. If you are having problems, check to see if you have an environment setting like the following on your local system:

```
_JAVA_OPTIONS=-Duser.home=...
```

When you install the command-line interface software, the default profile is installed in the profile directory with the software. The file name is `dscli.profile`. For example. `c:\Program Files\IBM\DSCLI\profile\dscli.profile`.

Example profile file modification: The following steps represent a typical modification process for some of the key items. Making these changes before you begin using DS CLI commands could save you time.

1. Click the DSCLI icon on your desktop (a DOS window opens)
2. Enter `edit dscli.profile` at the command prompt. The profile configuration file is displayed.
3. Scroll down to the # sign in front of HMC1: and remove the # sign.
4. Enter the correct IP address of your management console.
5. Scroll down to the # sign in front of DEVID and remove the # sign.
6. Enter the serial number of your machine type (include the values for manufacture, machine type, and serial number).
7. Save the file.
8. Enter `cd..` at your command prompt.
9. Enter DSCLI at your command prompt and the DS CLI applications starts. You are asked to provide only your user ID and password and not the address of your management consoles.

Table 3 provides the list of profile variables that you can use to create the profile.

Table 3. Profile variables

Variable	Description
banner: on off	Enables or disables the banner that appears before the command output. This variable is equivalent to the command option -bnr . The command option -bnr overrides this default value.
delim	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	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.
format	Specifies the output format for list commands. Specify one of the following formats: <ul style="list-style-type: none"> • 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 horizontal table. This variable is equivalent to command option -fmt . The command option -fmt overrides this default value.
fullid	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	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	Specifies the secondary Storage Manager IP address. This variable is equivalent to the command option -hmc2 . The command option -hmc2 overrides this default value.

Table 3. Profile variables (continued)

Variable	Description
locale	Specifies the language for the output on the local computer. <ul style="list-style-type: none">• ar: Arabic• be: Byelorussian• bg: Bulgarian• ca: Catalan• cs: Czech• da: Danish• de: German• el: Greek• en: English• es: Spanish• et: Estonian• fi: Finnish• fr: French• gu: Gujarati• hi: Hindi• hr: Croatian• hu: Hungarian• in: Indonesian• is: Icelandic• it: Italian• iw: Hebrew• ja: Japanese• kk: Kazakh• kn: Kannada• ko: Korean• lt: Lithuanian• lv: Latvian (Lettish)• mk: Macedonian• mr: Marathi• ms: Malay

Table 3. Profile variables (continued)

Variable	Description
locale, <i>continued</i>	<ul style="list-style-type: none"> • 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
paging: on off	Controls the display of output. If paging is enabled, a limited number of lines of output displays when a command is issued. The lines do not scroll. You must set the number of lines per page with the rows variable. This variable is equivalent to command option -p . The command option -p overrides this default value.
timeout	Set timeout value of client/server synchronous communication. The unit of the value is second. The default value is 420 seconds. You can set this timeout if the processing of a command ends by timeout due to network or client/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	Set 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: <ol style="list-style-type: none"> 1. If the DS CLI returns a connection error, check for the following conditions: <ul style="list-style-type: none"> • 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 too short can cause unexpected connection problems.

Table 3. Profile variables (continued)

Variable	Description
remotedevid	Specifies the remote storage image ID. This variable is equivalent to the command option -remotedev . The command option -remotedev overrides this default value.
rows	Specifies the number of rows per page of output if the paging variable is enabled. This variable is equivalent to command option -r . The command option -r overrides this default value.
verbose: on off	Enables or disables verbose output. This variable is equivalent to the command option -v . The command option -v overrides this default value.

Example

```
#
# DS CLI Profile
#

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

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

#
# 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 420 seconds.
#timeout: 900

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

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

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

```

paging: off
#rows: 24

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

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

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

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

# End of Profile

```

Removing the DS Command Line Interface

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

The following topics describe the steps required to successfully remove this interface.

Removing the DS CLI from your system using graphical mode

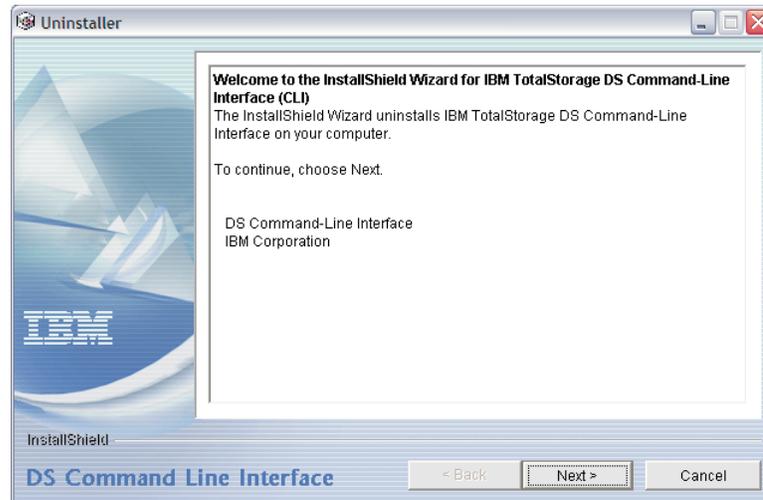
Use the graphical mode to remove the DS CLI from your system when the DS CLI is installed on a Windows, Novell, or UNIX system.

Notes:

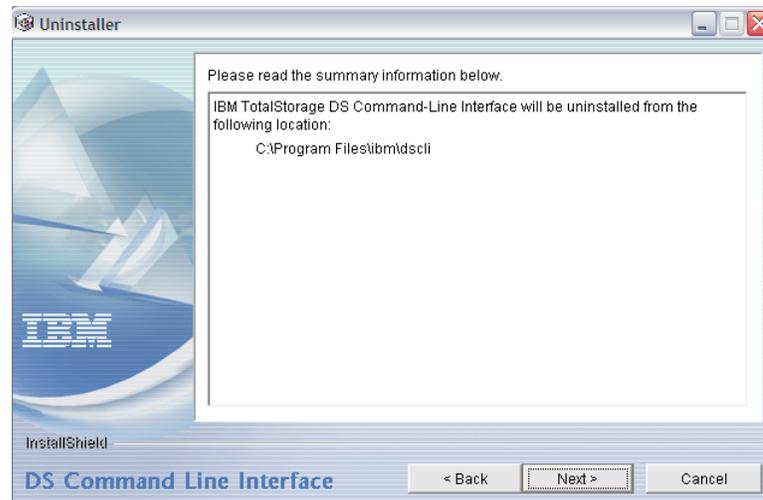
1. The following procedure applies to the removal of only the DS CLI application. This procedure cannot be used to remove other versions of the CLI.
2. If you do not want to create a new profile when you reinstall the CLI, select to not delete the DS CLI profile as you complete this task.

You use the Add/Remove Programs facility of the Windows operating system to remove the DS CLI from your system. When you have processed the uninstall steps, restart your system to complete the uninstall. Perform the following steps to remove the DS CLI using the graphical mode.

1. Navigate to the Windows Control Panel and open the Add/Remove program facility.
2. Scroll the list of currently installed programs and click the listing for IBM System Storage DS Command-Line Interface.
3. Click the **Change/Remove** button and the Welcome window for the Uninstaller is displayed.



4. 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 removed.



5. Click **Remove** to continue or **Cancel** to stop the removal and exit the uninstall process. Click **Back** to return to the previous window. When you click **Remove**, the Uninstallation Progress window is displayed. When the uninstall process is finished, the Finish window is displayed, which contains a statement about the success or failure of the uninstall process. Click **Finish** to close.



If the uninstall 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.

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

Removing the DS CLI using unattended (silent) mode

Use the unattended (silent) mode to remove the DS CLI through the command line if the DS CLI is installed on a UNIX system or a variant of UNIX (for example, Linux, HPUX, Sun, or AIX).

Perform the following steps to successfully uninstall the DS CLI.

Notes:

1. If you are using Windows or Novell, you will use the Add/Remove Programs feature to uninstall the DS CLI.
2. This uninstall process only works with DS CLI. No other versions of CLI can be removed with this process.
1. 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.xxx`, where `xxx` depends on the operating system. If you have a Linux, Hewlett Packard, Sun, or AIX system, then the file name is `uninstaller.bin`. For all other operating systems the file name is `uninstaller.sh`.
2. Type the following command at the command prompt: `<install directory>/_uninst/uninstaller.<exe | sh | bin> -silent`
3. Press the **Enter** key. All the associated CLI files are uninstalled.

Removing the DS CLI using the console mode

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

Perform the following steps to remove the DS CLI using the console mode:

Note: Do not use the console method to uninstall DS CLI on a Windows system. Instead, follow the steps in this guide for removing the DS CLI using graphical mode.

1. Type the following command at a command prompt: <install directory>/_uninst/uninstaller.<sh | bin> -console
2. The Welcome screen displays. Press 1 and Enter to continue, or 3 to Cancel the removal process.

```
Welcome to the InstallShield Wizard for IBM System Storage DS Command-Line
Interface (CLI)
The InstallShield Wizard uninstalls IBM System Storage DS Command-Line Interface
on your computer.

To continue, choose Next.

DS Command-Line Interface
IBM Corporation

Press 1 for Next, 3 to Cancel or 4 to Redisplay [1]
```

3. The Uninstallation Location screen is displayed. Press 1 and Enter to continue, or 3 and Cancel to exit the removal process.

```
IBM System Storage DS Command-Line Interface will be uninstalled from the
following location:

C:\Program Files\ibm\dsccli

Press 1 for Next, 2 for Previous, 3 to Cancel or 4 to Redisplay [1]
```

4. The Uninstallation progress screen is displayed while the command-line interface is being removed.

```
Uninstalling IBM System Storage DS Command-Line Interface...
```

5. The Uninstallation Finish screen is displayed. Press 3 to finish the removal.

```
The InstallShield Wizard has successfully uninstalled IBM System Storage DS
Command-Line Interface. Choose Finish to exit the wizard.

Press 3 to Finish or 4 to Redisplay [3]
```

Removing the DS CLI from a System i system

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

Because the DS CLI is installed on a System i system 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 system, you used a remote system to perform 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 system. 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 iSeries system.
You might use this method if you are not planning to upgrade the DS CLI and you want to totally remove the DS CLI from your System i system.
- 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.

Removing the DS CLI using your System i system directly

Complete this task to remove the DS CLI through the direct use of your System i system.

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 system 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 system for the removal process. However, it is possible to use your System i system directly (bypassing the uninstaller) to remove the DS CLI application.

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 system.

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

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.

Perform the following steps to remove the DS CLI through the direct use of your System i system:

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

```
RUNJAVA 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.

Removing the DS CLI from your System i system using the remote method

Complete this task to remove the DS CLI from your System i system 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 system 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 system for the removal process. However, it is possible to use the remote removal method on your System i system to remove the DS CLI application.

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

You can use the following remote method to remove the DS CLI from a System i system. You can perform this procedure at any time. However, it is common to perform this procedure when you want to upgrade the DS CLI application, 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 system using the remote method, perform the following steps:

1. Use FTP to copy the `uninstall.dat` and `uninstall.jar` files from the `uninstall` directory on your System i system into a common directory on the machine that you are using to do the uninstall.
2. Invoke the uninstall process from the common directory (or point to it using the `classpath [-cp]` option) by issuing the following command: `java -cp uninstall.jar run -os400 <system> <userid> <password>`
3. Wait until the uninstall process is completed before you continue to your next process.

OpenVMS system integration

You can adjust your OpenVMS system to obtain greater benefits from the use of the DS CLI application. 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 `1GGAn`. In contrast, fibre-channel disk devices have the OpenVMS device type `DG`, which result in device names in the form `1DGAn`. 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, and 1750 volume IDs. That implies that all host devices belonging to 2105/2107/1750 volumes are displayed. Therefore, CCL devices `1GGAn` 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 `1GGAn` CCL devices are shown in the command output. Otherwise, the `1GGAn` 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 `1GGAn` 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 **IBMDCLI\$MANAGER:IBMDCLI\$LOGIN.COM** activates the message section file **IBMDCLI\$SYSTEM:IBMDCLI_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 **IBMDCLI\$HELP**:

IBMDCLI_Ovr.hlp

A help library containing one module with the top-level key **IBMDCLI**. 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 forth.

IBMDCLI_Messages.msghlp\$data

A message help data file with messages for facility **IBMDCLI**. 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 **IBMDCLI\$MANAGER:IBMDCLI\$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 **IBMDCLI\$JRE**. This setup is according to HP guidelines. The login procedure **IBMDCLI\$MANAGER:IBMDCLI\$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 **IBMDCLI\$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:

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

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

Chapter 4. Configuring and managing logical storage

This section provides the information that you need to configure and maintain your logical storage whether it consists of fixed block volumes or count-key-data volumes.

You must complete a logical configuration for each storage unit. For example, each storage unit 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 using the DS CLI commands.

Configuring the DS8000 (using DS CLI) for use with the TotalStorage Productivity Center Replication Manager

Complete this task to configure the DS8000 for use with the TotalStorage Productivity Center Replication Manager using the DS command-line interface.

The following items must be in place before you can configure the DS8000 for TotalStorage Productivity Center Replication Manager usage:

- Ethernet cards (2 per storage image) must be installed. The Ethernet cards are 2-port cards. Only the lower-numbered port, which is the top port on the card, is connected to the external IP network. It is this port that must be configured. **Do not configure the 2nd port on the card.**
- TotalStorage Productivity Center Replication Manager V3.1.2 or later must be installed on a separate server.
- The DS command-line interface must be installed.

TotalStorage Productivity Center Replication Manager is an optional software feature that can be used to manage replication for the DS8000. TotalStorage Productivity Center Replication Manager allows you to easily manage large volume Copy Services functions. Activation of this feature for the DS8000 requires the configuration of the DS8000 using the DS command-line interface commands. It also requires a configuration of the TotalStorage Productivity Center Replication Manager using the TotalStorage Productivity Center Replication Manager GUI or CLI.

You must establish a communication link between the DS8000 and the server that contains the IBM TotalStorage Productivity Center for Replication software program. To do this requires that you obtain the following key information from your network administrator:

- TCP/IP address for the designated port
- TCP/IP subnet mask for the designated port
- Gateway TCP/IP address for the designated port
- Primary DNS TCP/IP address for the designated port
- Secondary DNS TCP/IP address for the designated port

Note: The User ID and password that you use to allow the IBM TotalStorage Productivity Center for Replication software program to access the DS8000 are as follows:

User ID - *tpcruser*

Password - image ID (for example, 75BZFA1). This is not the unit number. The letters within the password value must be entered as caps otherwise the login fails.

You can configure your Ethernet cards at any time. However, you cannot test the connection until at least one LSS or one LCU has been configured on your DS8000.

Perform the following steps to configure your Ethernet cards and establish a connection with the IBM TotalStorage Productivity Center for Replication software program:

1. Obtain the key information from your network administrator.
2. Log in to the DS CLI application in interactive mode.
3. Follow the configuration steps to create at least one LSS or one LCU.

Note: You can skip this step if storage is already configured on your DS8000.

The configuration steps include the following tasks:

- a. Create extent pools
- b. Create arrays
- c. Create a rank
- d. Create a logical control unit (CKD only)
- e. Create volumes
- f. Create volume groups (Open systems only)
- g. Configure your fibre-channel I/O ports
- h. Create the SCSI host port connections (Open systems only)

Note: All of these steps are considered part of the initial configuration process and they are explained in the DS8000 Information Center command-line interface topic *Configuring and managing logical storage*.

4. Issue the **lsnetworkport** command using the **-l** parameter to generate a report that lists the different network ports that are available in the specified storage unit. In addition, the report displays the Port ID values for the ports that you will configure using the **setnetworkport** command.

The following types of information are displayed on this report:

State Specifies whether a port is enabled or in some other state.

Location

Specifies the network port location using the format:

Utttt.mmm.ppssss-Pn-Cn-Tn, where

- *Utttt.mmm.ppssss* identifies the location of the processor complex.
- *Pn* identifies the planer number.
- *Cn* identifies the card number.
- *Tn* identifies the port number.

5. Issue the **setnetworkport** command to configure the Ethernet cards for the DS8000 after you have configured at least one LSS. Enter the **setnetworkport** command at the dscli command prompt with the following parameters and variables:

```
dscli>setnetworkport -dev storage_image_id -ipaddr IP_address  
-subnet IP_mask -gateway IP_address  
-primary IP_address -secondary IP_address port_ID
```

Example

```
dscli>setnetworkport -dev IBM.2107-75FA120 -ipaddr 9.11.555.555  
-subnet 255.555.555.0 -gateway 9.11.555.0 -primary 5.5.5.5 -secondary 5.0.5.0  
19801
```

Notes:

- a. Only those ports whose ID ends with a value of 01 (for example, 19801) should be configured using the **setnetworkport** command. Do not attempt to configure the ports that end with a value of 02. (The list of port IDs was provided in the report generated by the **lsnetworkport -l** command.)
 - b. Values for the **-ipaddr**, **-subnet**, and **-gateway** parameters are required.
 - c. You can optionally configure the primary and secondary DNS addresses.
 - d. The specified IP addresses do not respond to network ping commands.
 - e. You must process the **setnetworkport** command once for each ethernet adapter that is installed. This equates to entering the command twice for a single SFI machine, or either twice or four times for a dual-SFI machine (depending on whether the cards are installed for both SFIs or not).
6. Set up the communication link on the server that contains the IBM TotalStorage Productivity Center for Replication software program.
- a. Log in to TotalStorage Productivity Center Replication Manager program and follow the menu path **Storage Subsystem** → **Add Subsystem** → **ESS/DS**
 - b. Fill in the applicable information which consists of the following:
 - IP addresses for the two Ethernet cards.
 - User ID (*tpcruser*) and password (storage image unit ID) that allow the TotalStorage Productivity Center Replication Manager program to access the DS8000.

Notes:

- 1) The password is the image ID of the storage unit not the unit number (for example, you might have a unit number of 75BZFA0 and the associated image ID is 75BZFA1. The image ID is the password.)
- 2) You can use the **lssi** command to verify the value that you must use for the password. The following example shows the report that is generated when you type **lssi** at the dscli command line:

Note: The table format displayed for the report is not how the actual report is displayed. The format is simply used for clarity.

```
dscli> lssi
```

```
Date/Time: September 15, 2006 10:39:32 AM MST DS CLI Version:  
5.2.200.404
```

Name	ID	Storage Unit	Model	WWNN
-	IBM.2107- 1301371	IBM.2107- 1301370	921	5005076303 FFC04E

State	ESSNet
Online	Enabled

The value that is highlighted in the above report (**1301371**) is the password that is used on the specified server.

- 3) Ensure that all alphabetic characters that are a part of your password are entered as caps otherwise the login will fail.

Click **OK** to process the transaction. The Storage Subsystems panel is displayed.

7. Verify that the connection link has been established between the DS8000 and the server that contains the IBM TotalStorage Productivity Center for Replication software program.

The Local Server connection column displays a green checked dot and the word "Connected" when the connection has been made.

Configuring new fixed block storage using the DS CLI

This section describes how you can configure new fixed block storage within a storage image by using the command-line interface.

Before you begin, you must be logged into the DS CLI application in interactive command mode. You must also be connected to a storage image 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 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, thereby 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 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 **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 in order 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 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 (10⁹ bytes). 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 application and connected to the storage unit that will be used for open systems host system storage.

Perform 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

1. Issue the **mkextpool** command to create the fixed block extent pool for rank group 0. Enter the **mkextpool** command at the dscli command prompt with the following parameters and variables:

```
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:

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
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 in order 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 **-l** 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:

```
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 machine type 2107 storage image storage devices (DDMs) are packaged into storage enclosure pairs. The machine type 2107 contains at least one storage enclosure pair, with a minimum of 16 DDMs.

The DDMs of a storage enclosure are partitioned into array sites. A machine type 2107 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. All 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 array site. An array inherits the characteristics of its parent array sites and is given a RAID type attribute (5 or 10). A 2107 array of RAID type 5 or 10 is made from one (8 DDM) array site. The status of the array is “unassigned” until the array is assigned to a rank.

Perform the following steps to create an array from unassigned array sites:

1. Issue 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. 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 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. Issue the **mkarray** command to create an array from one array 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 | 10] -arsite array_site
```

Consider the following when you create the arrays:

- 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.
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 application and connected to the storage unit that will be used for open systems host system storage.

To create ranks, perform the following steps:

1. Ensure you have a list of the unassigned arrays for which ranks must be assigned. Issue the **lsarray** command to obtain this list if you do not already have it. Enter the **lsarray** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsarray -dev IBM.2107-75FA120 -state unassigned
```

2. Issue 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. Enter the **mkrank** command at the dscli command prompt with the following parameters and variables:

```
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, you can issue the **lsrank** 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. Issue the **lsrank** command to verify that ranks and extent pools have been assigned. Enter the **lsrank** command at the dscli command prompt with the following parameters and variables:

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

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.

You must have completed the following tasks before you can create your fixed block volumes:

- Created your extent pools.
- Created your arrays.
- Created and assigned your ranks.

Perform 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. Enter the **lsextpool** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsextpool -dev IBM.2107-13AAD7A -stgtype fb -1
```

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. Issue the **lsaddressgrp** command to find unassigned and available address groups. Enter the **lsaddressgrp** command at the dscli command prompt with the following parameters and variables:

```
dscli> lsaddressgrp -dev IBM.2107-75FA120 -1
```

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 this is your first time creating fixed block volumes, 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. Issue the **mkfbvol** command to create fixed block 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 P1  
-name finance#d -cap 8.6 0100-010F
```

Consider the following conditions with regard to the command example in this step:

- All volumes will 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 wild card (#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 GB. The default is binary GB where 1 GB = 1 073 741 824 (2³⁰ 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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE, and *ZZ* is the volume number (00 - FF).
 - 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.
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: It is possible that the report will display that there was a configuration error that is associated with one or more of your **mkfbvol** transactions.

Creating LUN volumes for System i systems

Complete this task to create fixed block LUN volumes for System i systems.

You must have completed the following tasks before you can create your fixed block LUN volumes:

- Created your extent pools
- Created your arrays
- Created and assigned your ranks

When you begin your initial configuration, the LSSs and address groups do not exist. The LSS's are created when the first volume of the LSS's is defined during the processing of the **mkfbvol** command; likewise, 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 two 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 number FF is reserved for internal use and must not be used as a volume ID.
- You must define each volume as protected or unprotected. This is simply a notification to i5/OS; it does not mean that the volume is protected or unprotected. In reality, all LUNs are protected, either by RAID5 or RAID10. 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 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.

Perform 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, issue 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 -1
```

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. 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 with regard to the command example in this step:

- The **-extpool** parameter identifies a fixed block extent pool containing available data extents.
- The **-os400** parameter allows you to designate 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.
- 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 hexadecimal 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 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 4 hexadecimal digits in the form of XYZZ where X is the address group (0 - 1), XY together is the logical subsystem number 00 - FE, and ZZ is the volume number (00 - FF).

- You can define up to 255 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 number FF is reserved.
3. Repeat step 2 for all of the required 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 -1
```

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 have created.

The **lsfbvol** or the **showfbvol** commands display reports that includes a *configstate* category. The configuration state category reports on the current state of the rank. 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. Generally, to correct this error state, you must delete the designated volume configuration and submit a new transaction request.

Perform the following steps to obtain additional information about the configuration error and to correct this error condition.

1. Add the **-v** (verbose) command flag 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 rank 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 the majority of instances, this is the only method 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) at a later time. 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 will use the volume group.

Perform 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. Because you need to know the criteria that is associated with these volume group types, see the **mkvolgrp** command for more information.
 - 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 Os400mask** 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 2107.
 - The volume group name (*my_nickname* in the example command) must be unique within the scope of the specified storage image.
2. 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 -l
```

Notes:

- a. The **lsvolgrp** command with the **-l** parameter displays a report with the following 3 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 systems

Complete this task to create volume groups for System i systems 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.
- If you do not use an external load source, you create a volume group that contains all the volumes that are assigned to an i5 fibre-channel adapter. You cannot specify more than 32 LUN volumes for each attachment to an i5 fibre-channel adapter.

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, also you denote that the logical size of the blocks in the volumes is 520 bytes.

Perform the following steps to create and view volume groups for System i systems:

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
```

```
Date/Time: July 5, 2005 11:57:50 PM GMT+01:00 IBM DSCLI Version: 5.0.4.32
DS: IBM.2107-13ABVDA
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

Date/Time: July 5, 2005 11:57:50 PM GMT+01:00 IBM DSCLI Version: 5.0.4.32
DS: IBM.2107-13ABVDA
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. System i only supports 32 device addresses per volume group.
2. 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 -l
```

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:

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 will 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 systems 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 machine type 2107, HA cards are installed in I/O enclosures, each containing up to four HA cards. 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, perform 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 2107 machine type supports the “identified” access mode for SCSI host attachments, which requires that all SCSI host ports be identified to a storage image. One SCSI host port connection must be created for each SCSI host port that accesses storage image 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 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 **lspportprof** 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, perform the following steps:

1. 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:

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)

Name	Profile	AddrDiscovery	LBS	Description
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 -wwname wwpn
             -hosttype host_type -volgrp volume_group_ID -ioport port_ID
             host_name
```

Notes:

- a. The **-wwname** 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 **-I** parameter.

Configuring new count key data storage using DS CLI

This section describes how to configure new count key data (CKD) storage within your storage image for a zSeries host system.

Before you begin, you must be logged into the DS CLI application in interactive command mode. You must also be connected to a storage unit that is used for zSeries systems 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 zSeries 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) 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 that are configured for FICON I/O operations. 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

Complete this task to create CKD volume extent pools. This is the first step in configuring new count key data storage. You can use the DS CLI commands to create extent pools for CKD volumes.

Creating the extent pools before the arrays and ranks saves a processing step. When you create the 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**. At a minimum, 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 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. It is better to evenly distribute rank and extent pool allocations in order to keep the storage server workloads balanced.

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 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 as ranks.

Use the **lsextpool** and **mkextpool** commands to create the CKD extent pools. You must be logged into the DS CLI application and connected to the storage unit that will be used for your zSeries host system.

To make your extent pools, perform the following steps:

1. **Find defined CKD extent pools** by issuing the **lsextpool** command to display a list of the existing CKD extent pools. Type the following command at the dscli command prompt:

```
dscli> lsextpool -dev IBM.2107-75FA120 -stgtype ckd
```

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 this is the first time that the extent pools are created, the minimum number of extent pools does not already exist.

- Does each extent pool have a rank group that is assigned to it and are they balanced? If this is the first time, there will be no rank assignments.
- Are additional extent pools needed?

3. **Make your extent pools** by issuing the **mkextpool** command. (A minimum of two extent pools must be created. One for rank group 0 and one for rank group 1.)

Type the following command at the dscli command prompt:

```
dscli>mkextpool -dev IBM.2107-75FA120 -rankgrp 0 -stgtype ckd  
extent_pool_name  
dscli>mkextpool -dev IBM.2107-75FA120 -rankgrp 1 -stgtype ckd  
extent_pool_name
```

The *extent_pool_name* parameter is required with the **mkextpool** command. The extent pool name is a unique name that you assign to each extent pool and it cannot be longer than 16 characters.

Make additional extent pools for each of the following conditions:

- Each RAID type (5 or 10)
- Each disk drive module (DDM) size
- Each CKD volume type (3380, 3390)
- Each logical control unit (LCU) address group

4. **Press Enter.** A successful process displays the following message:

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.

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
Extent pool P1 successfully created.
```

5. **Repeat Step 2** for each extent pool that you want to create. Try to evenly distribute rank and extent pool allocations in order to keep the storage server workloads balanced.
6. **Verify the extent pool assignments** by issuing the **lsextpool** command when you are done making the extent pools. Use the **-l** parameter to display a full

report for the extent pools that are assigned to the storage unit. Type the following command at the dscli command prompt:

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

You might want to print this report because this list is used later during the process of creating 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 machine type 2107 storage image storage devices (DDMs) are packaged into storage enclosure pairs. The machine type 2107 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 DDMs of a storage enclosure are partitioned into array sites. A machine type 2107 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. 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 or 10). A 2107 Array object of RAID type 5 or 10 is made from one (8 DDMs) array site.

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 application and connected to the storage unit that will be used for open systems host system storage.

Perform the following steps to create arrays for a CKD volume configuration:

1. 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 | 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 RAID types 5 and 10. 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. 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 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 application and be connected to the storage unit that will be used for open systems host system storage.

To make ranks, perform the following steps:

1. Issue the **lsarray** command to ensure you have a list of the unassigned arrays for which ranks must be assigned. Enter the **lsarray** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsarray -dev IBM.2107-75FA120 -state unassigned
```

2. Issue 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. Enter the **mkrank** command at the dscli command prompt with the following parameters and variables:

```
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 allows you to be notified if the rank configuration has failed for any reason.
 - b. Be aware that when you use the **-wait** parameter, you cannot issue any other commands until the entire transaction has processed.
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, issue the **lsrank** command from a different DS CLI session. A successful process generates 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 and step 3 until all unassigned arrays are assigned a rank and an extent pool.
5. Issue the **lsrank** command to verify that ranks and extent pools have been assigned. Enter the **lsrank** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsrank -dev IBM.2107-75FA120 -l
```
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 S/390 and zSeries 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.

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 application and connected to the storage unit that will be used for open systems host system storage.

To create LCUs, perform the following steps:

1. 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:

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.

lcu type

Specifies the type of LCU to be created. You can specify the following types:

- 3390-3
 - 3990-tp
 - 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 **-l** 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).

Perform 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 wild card (#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 rank. 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. Generally, to correct this error state, you must delete the designated volume configuration and submit a new transaction request.

Perform the following steps to obtain additional information about the configuration error and to correct this error condition.

1. Add the **-v** (verbose) command flag 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 will 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 setiport 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 setiport 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 systems using i5/OS level V5R3M5 and above.

ficon (coded as ficon in the setiport 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 machine type 2107, HA cards are installed in I/O enclosures, each containing up to four HA cards. Use the **lsioport** and **setiport** 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, perform 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>setiport -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

The topics in this section provide task information that are related to working with configuration files after the initial creation of the configuration. This information is for working with the commands associated with the DS CLI application.

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 performing 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.
- You have configured the DS CLI profile as appropriate. To edit the profile file, perform the following:
 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 performed an IPL to System i™.

Assuming that you have met the previous conditions, perform the following steps using DS CLI on i5/OS to process storage configuration and Copy Services functions on the DS8000:

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 Menu                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 Assistant
F23=Set initial menu
```

2. **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

HMC1 . . . . . *PROFILE
HMC2 . . . . . *PROFILE
User . . . . . 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.

```
Date/Time: July 8, 2005 2:55:20 PM MST IBM DSCLI Version: 5.0.4.32 DS:
IBM.1750-13ABVDA
```

```
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 application.

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

Perform the following steps to modify the extent pool properties.

1. Issue the **lsxtpool** command to generate a report that identifies the status of your extent pools by storage type (fixed block or count key data). Enter the **lsxtpool** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsxtpool -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 **lsxtpool** 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 unit and status information on each extent pool in the list.

Issue 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 unit. 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, issue the **showextpool** command.

Perform the following step to display a list of extent pools and their status in a storage unit.

Issue the **lsextpool** command to display the extent pool list and status information. Enter the **lsextpool** command at the dscli command line prompt with the following parameters and variables:

```
dscli>lsextpool -dev storage_image_ID -l
```

Notes:

1. Use the **-l** parameter if you want to see the list and status for all the extent pools (fixed block and CKD) in your storage unit. A full report is displayed.
2. Use the **-s** parameter if you just want to see only a list of the extent pools in your storage unit. 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 unit. You can obtain these IDs by issuing the **lsextpool** command.

Issue 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 unit.

Perform the following steps to display the detailed properties of an extent pool or to display the performance metrics of an extent pool.

1. (For detailed properties information) Issue the **showextpool** command. Enter the **showextpool** command at the dscli command-line prompt with the following parameters and variables:

```
dscli>showextpool -dev storage_image_ID extentpool_ID
```
2. (For performance metrics information) Issue the **showextpool** command. Enter the **showextpool** at the dscli command-line prompt with the following parameters and variables:

```
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 unit 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 **lsxtpool** command. After you determine which extent pools can be deleted, you can issue the **rmxtpool** command that designates the extent pools that you want to delete.

Perform the following steps to delete one or more extent pools from a fixed block volume configuration.

1. Issue the **lsxtpool** command to display a list of extent pools. Ensure that you designate the storage type within your command parameters. Enter the **lsxtpool** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsxtpool -dev storage_image_ID -l -stgtype [fb | ckd]
```

2. Analyze the list and determine which extent pools can be deleted.
3. Issue the **rmxtpool** command to delete the designated extent pools. Enter the **rmxtpool** command at the dscli command prompt with the following parameters and variables:

```
dscli>rmxtpool -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 **lsxtpool** 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 machine type 2107 storage image DDMs are packaged into storage enclosure pairs. The machine type 2107 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 DDMs of a storage enclosure are partitioned into array sites. A machine type 2107 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. 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.

Perform the following steps to view the DDM status:

1. Issue the **lsddm** command to obtain a list and status of the DDMs currently associated with the storage unit. Enter the **lsddm** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsddm -l storage_image_ID
```

2. 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 need to 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 unit. 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.

Perform the following steps to view the status of all the array sites and arrays that are associated with the storage unit.

1. Issue the **lsarraysite** command to generate a list of all the array sites and their status. Enter the **lsarraysite** command at the dscli command prompt with the following parameters and variables:

```
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:

- Assigned - The array site has been defined as an array.
 - Unassigned - The array site is available to be defined as an array.
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 -l
```

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 unit using DS CLI commands.

The steps in this task presume that you have created your arrays.

The creation of arrays is based on the array sites that are associated with the storage unit. After you have created your arrays, there might be times when you want to view the status of an array site and the associated array.

Perform the following steps to view the status of the array site or sites and the array that is associated with the storage unit.

1. Issue the **showarraysite** command to generate a report that displays the array site or sites and their status. Enter the **showarraysite** command at the dscli command prompt with the following parameters and variables:

```
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. The array site state shows a value of **assigned** if you have created the array.
2. Issue the **showarray** command to generate the properties report for the specified array. Enter the **showarray** command at the dscli command prompt with the following parameters and variables:

```
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.

Perform 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 generally 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.

Perform the following steps to add a rank to an extent pool:

1. Issue the **lsrank** command to generate a report that lists the status of the ranks that are associated with the storage unit. Enter the **lsrank** 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 **chrnk** command to add (reassign) a rank or ranks to an extent pool. Enter the **chrnk** command at the dscli command prompt with the following parameters and variables:

```
dscli>chrnk -dev storage_image_ID -extpool extentpool_ID rank_ID
```

Notes:

- a. The rank ID is a 4-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 perform 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.

Perform the following steps to modify a rank using DS CLI commands:

1. Issue the **lsrank** command to generate a report that lists the status of the ranks that are associated with the storage unit. Enter the **lsrank** 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. In other words, 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 unit 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 unit. 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

Perform the following step to view the status of all the ranks that are associated with the storage unit.

Issue the **lsrank** command to generate a list of all the ranks and their status. Enter the **lsrank** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsrank -dev storage_image_ID -l
```

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 unit 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.

Perform the following steps to view the status of a rank that is associated with the storage unit.

Issue the **showrank** command to generate the properties report for the specified rank. Enter the **showrank** command at the dscli command prompt with the following parameters and variables:

```
dscli>showrank -dev storage_image_ID rank_ID
```

Notes:

1. Because the **showrank** command requires the use of a specific rank ID, you can issue 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. Generally, to correct this error state, you must delete the designated rank configuration and submit a new transaction request.

Perform 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 flag 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.

Perform the following steps to remove ranks from a storage configuration:

1. Issue the **lsrank** command to obtain a list of the ranks that are associated with the storage configuration that is being deleted. Enter the **lsrank** 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 perform 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

Perform 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>chlsu -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 148.
 - d. To modify the consistency group timeout value, see “Modifying the consistency group timeout value” on page 148.
 - e. To modify the critical mode (Administrator authority only), see “Modifying the critical mode setting” on page 149.
 - f. To modify the z/OS Global Mirror timeout value, see “Modifying the z/OS Global Mirror timeout value” on page 149.

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 unit 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 unit. 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

Perform the following step to view the status of all the LCUs that are associated with the storage unit.

Issue the **lslcu** command to generate a list of all the LCUs and their status. Enter the **lslcu** command at the dscli command prompt with the following parameters and variables:

```
dscli>lslcu -dev storage_image_ID -l
```

Notes:

1. Issue 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 unit using DS CLI commands.

An LCU represents a logical subsystem for zSeries 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.

Because you can modify some of the properties of an LCU, there might times when you want to examine the associated properties. The **showlcu** command allows you to view the properties of a single LCU.

Perform the following step to view the properties of a single LCU:

Issue the **showlcu** command to view a report that displays the properties of a single LCU. Enter the **showlcu** command at the dscli command prompt with the following parameters and variables:

```
dscli>showlcu -dev storage_image_ID LCU_ID
```

Note: The LCU ID is a 2-digit hexadecimal number in the range of 00 - FE .

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

Perform the following steps to remove LCUs from a CKD storage configuration:

1. 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).
3. 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 5. Copy Services functions

This topic provides information about how to use DS CLI commands to perform 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 the DS8000 storage units and are supported on many operating systems. For example, if you set up and configure your DS8000 to use i5/OS, you can create copies of System i™ disk pools within a single DS8000 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 actual physical copy processing takes can depend on the amount of data being copied and other activities that are occurring on the storage unit.

You can perform this task using either the DS CLI or the DS Storage Manager.

Perform 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. Issue the **mkflash** command to create FlashCopy relationships. 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
```

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. 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 **-l** 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 perform 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 perform this task using either the DS CLI or the DS Storage Manager.

Perform 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” on page 501 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

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
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 perform this task using either the DS CLI or the DS Storage Manager.

Perform 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. A separate example is not shown for the 1750 because the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

ID	SrcLSS	Sequence Num	Time-out (secs)	Active Copy	Recording	Persistent	Reversible
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	Background-Copy	Copy-Indicator	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 perform this task using either the DS CLI or the DS Storage Manager.

Perform 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.

1. 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 **-l** 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 perform 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 similar to the following 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 DS Storage Manager. (Part of the Remote FlashCopy operation requires that you create paths and volume pairs first. You can issue those requests using either the DS Storage Manager 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 DS Storage Manager, because network connections to both the source and target sites are required. If the network connection to the target site is lost, the DS Storage Manager cannot connect to the target site. Whether you use the DS Storage Manager 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 DS Storage Manager for this step because connections to the target site are lost.

Assume that you performed Step 1 and Step 2 from the DS Storage Manager (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.

Date/Time: February 7, 2005 4:29:42 PM IST IBM DSCLI Version: 5.0.0.0 DS: IBM.2107-1300871 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 performed 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 new 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 performed are copied to the target volume. The specified parameters in this command replace the parameters in the existing relationship. In order to keep the initial **-record** and **-persist** parameters, specify the **-record** and **-persist** parameters with the **resyncflash** command.

You can perform this task using either the DS CLI or the DS Storage Manager.

Perform 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

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
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 perform this task using either the DS CLI or the DS Storage Manager.

Perform 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

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

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

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 back 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 prior to performing this task. It is not valid to perform 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 performed. Both the commit changes and discard changes tasks disable the FlashCopy revertible option.

You can perform the FlashCopy Revertible task using either the DS CLI or the DS Storage Manager.

Perform 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: 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 dsccli command prompt with the following parameters and variables:

```
dsccli>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

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
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.

Perform 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 **-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.
```

Example

```
dscli>lsflash -dev IBM.2107-75FA150 -l 0100:0100
```

Use the **-l** 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 perform this task using either the DS CLI or the DS Storage Manager.

Perform 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 perform this task using either the DS CLI or the DS Storage Manager.

Perform 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. In order to keep the initial **record**, **persist**, and **tgtpprc** parameters, the **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

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

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 perform this task using either the DS CLI or the DS Storage Manager.

Perform 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 performed 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 new 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 perform this task using either the DS CLI or the DS Storage Manager.

Perform 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 performed 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

This topic provides information to help you use Metro Mirror functions using the DS CLI commands. 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.

Note: If you are using the Cisco MDS 9216 Multilayer Fabric Switch, you must not enable its write acceleration feature. The Remote Mirror and Copy 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.

Perform 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 performed.

- 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 **lsavailpprport** and **mkpprpath** commands.

Perform 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 **lssi** command to display the list of WWNNs. Enter the **lssi** 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 performed. 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 that will 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. These are the paths 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.

Perform 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 **mkpprpath** command to create the fibre-channel paths for the remote mirror and copy source and target volume pairs. Enter the **mkpprpath** command at the dscli command prompt with the following parameters and variables:

```
dscli>mkpprpath -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 -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 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. 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.

Perform 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 **shows** command to obtain the remote WWNN number of the secondary storage unit.

tgflss 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.

Perform 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 similar to the following 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.

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 be performed. 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 storage unit and are supported on many operating systems. For example, if you set up and configure your DS8000 to use i5/OS, you can use Metro Mirror to create a copy of a System i disk pool on a separate DS8000, typically in a remote location.

Perform 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_imag_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 be performed.

In order 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 using consistency groups. A consistency group is a group of volumes that provides the ability to temporarily queue (at the host's 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.

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 hold off I/O activity through the **freezepprc** command. The unfreeze operation allows I/O activities to resume when you issue the **unfreezepprc** command.

You can specify the consistency group option when you are defining Metro Mirror paths between pairs of LSSs or 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. In addition, the group of

volumes in a consistency group can also consist of source volumes that are associated with a Model 2107 storage unit and target volumes that are associated with an ESS 2105 Model 800 or 750.

Perform 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.

1. 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
```

2. View the current consistency group state setting status of the consistency group by issuing the **showlss** command. You can also use the **chlss** command to change the default consistency group timeout value.

```
dscli>showlss -dev storage_image_ID LSS_ID
```

Example

```
dscli>showlss IBM.2107-75FA120/10
```

Resuming a Metro Mirror relationship

Complete this task to resume a Metro Mirror volume pair that has been suspended (paused).

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 performed. If you are using a Model 2105 ESS as part of the configuration, ensure that you have enabled the PPRC Version 2 license.
- 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.

Perform 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.

Issue the **resumepprc** command to continue Metro Mirror processing after it has been suspended (paused). Enter the **resumepprc** command at the dscli command prompt with the following parameters and variables:

```
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

Example

```
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 perform 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.

Perform 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 be performed. 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 storage units and are supported on many operating systems. For example, if you set up and configure your DS8000 to use i5/OS, you can use Global Copy to create a copy of a System i disk pool on a separate DS8000, typically in a remote location.

Perform 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 **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 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.

Perform 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. Issue 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. 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 -l -remotedev IBM.2107-75FA150 0100:0100  
0101:0101
```

2. Issue the **rmpprc** command to delete the Metro Mirror relationship between the source and target volume. Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

```
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 has 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 **mkpprpath** 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 zSeries 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.

Perform 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.

Perform 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.

Perform 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 perform 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.

Perform 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 have to be IPLed in order 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 zSeries 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	<ul style="list-style-type: none"> Suspends the primary volume. Allows write operations to the primary volume.
Critical Volume	Disabled	Enabled	<ul style="list-style-type: none"> Suspends primary volume when the last path to the secondary volume has failed. Inhibits write operations to the primary volume.
Critical Heavy	Enabled	Enabled	<ul style="list-style-type: none"> 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 be performed.

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, 1750:1750, 1750:2105.

A consistency group is a group of volumes that provides the ability to temporarily queue (at the host's 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.

Perform 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.

Perform 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
```

Performing a failback recovery operation

Complete this task to perform 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 in order 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 in order to resume mirroring.

Perform the following step to perform a failback 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.

1. Issue the **failbackpprc** command to perform a failback recovery 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 **lspprc** command to check the status information of each Metro Mirror relationship in the list. 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 -s  
SourceVolumeID:TargetVolumeID
```

Example

```
dscli>lspprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150  
-s 0100:0100 0200:0200 0300:0300
```

Performing a failover recovery operation

Complete this task to perform a failover recovery 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.

You must have created 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 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.

Perform the following step to perform 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 perform a failover recovery 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, and status information for each remote mirror and copy volume relationship in the list.

Perform 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 -l -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 Model 2107 storage unit and the target volume was associated with an ESS 2105 Model 800 or 750, you can convert that volume pair to synchronous.

Perform 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 performed.
- 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 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.

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.

Perform 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 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. For example, your command line would look like the following:

```
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 visa versa. If this 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 has failed.

To correct this situation and to remove the other part of the pair relationship, you must reissue the **rmpprc** command for each volume that was not removed using the following parameters:

- The **-at src** parameter, if the pair relationship has not been removed from the source volumes.
- The **-at tgt** parameter, if the pair relationship has not been 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.

Perform 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. Issue the **lspprc** command to generate a report about Metro Mirror relationships. This can help you determine which Metro Mirror relationships must be deleted. 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 -l 0100:0100 0101:0101
```

2. Issue the **rmpprc** command to delete the pair relationship that the volume still maintains. Enter the **rmpprc** command at the dscli command prompt with the following parameters and variables:

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 TargetVolumeID
```

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 2107/242x

Copy Services functions that are performed using either the DS Storage Manager or DS CLI interact with both the DS8000 (Machine type 2107/242x) and the IBM System Storage 2105 Enterprise Storage Servers (ESS) Models 800 and ESS 750.

Most Copy Services functions that are available on the ESS 2105 are also available on the 2107/242x and 1750 and in open systems and zSeries environments.

On the 2107/242x, 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 2107/242x machine which can be seen by the 2105 machine. You can use the mirroring solutions that are compatible between the 2105 and 2107/242x 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 perform Copy Services functions between machine types 2105 and 2107/242x, you must configure a Copy Services domain on the DS Storage Manager 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 performed 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 2107/242x. Otherwise, you must add them using either the DS Storage Manager 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 2107/242x, you must install licensed internal code version 2.4.2 or later on the 2105.
- The 2105 supports remote mirror and copy (formerly PPRC) operations with an ESCON link. However, if you want to configure a remote mirror and copy relationship between the 2107/242x and 2105, you must use a FCP link.

Creating a Metro Mirror volume pair between a 2107 and a 2105

Complete this task to create a Metro Mirror volume pair using volumes from a 2107 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 Model types 2107 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 2107 must be configured using Fibre Channel Protocol (FCP) ports.
- The storage type of the source and target volumes on the 2107 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 DS8000.

You can create Metro Mirror relationships using source and target volumes from the following machine types:

- A 2107 and a 2107
- A 1750 and a 1750
- A 2107 and a 1750
- A 2105 and a 1750
- A 2105 and a 2107
- A 2105 and a 2105

Note: If the source is a Copy Services 2105 domain, the Metro Mirror task is performed on the source domain. However, if you perform a "Suspend at target" action, the suspension occurs at the target domain.

Perform the following steps to create a Metro Mirror pair between a 2107 and a 2105. For this task, the source domain is a 2105 Model 800 or 750 and the target is a 2107. You can use this task if the target domain is a 2105 Model 800 or 750 and the source is a 2107 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.

1. Issue the **lsavailpprcport** command to list all of the available ports to the remote system to create a connection. It is recommended 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 to the ESS 800. You associate an LSS on the DS8000 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 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 **lspprcpath** 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 **lspprcpath** command at the dscli command prompt with the following parameters and variables: dscli> lspprcpath -dev IBM.1750-13AB7DA 10

Example

```
dscli>lspprcpath -dev storage_image_ID Source_LSS_ID
```

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.1750-13AB7DA -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.1750-13AB7DA -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.1750.13AB7DA 1001-10ff
```

Global Mirror functions

This topic provides information to help you use Global Mirror functions using the DS CLI commands. 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.

Note: If you are using the Cisco MDS 9216 Multilayer Fabric Switch, you must not enable its write acceleration feature. 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.

Remote flash copy pairs must be mapped to the target storage unit to create a valid global mirror.

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 volume which, for example, is converted from Global Copy to Metro Mirror, the formation of a 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.

Perform the following steps to add volumes to a Global Mirror session. 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 **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 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 -l 0100:0100
```

2. 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 parameters and variables shown as follows:

```
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 is displayed that indicates the session has been modified successfully.

3. Issue the **lssession** command to query the status of all volumes being processed including the volumes you added to the session. Enter the **lssession** command at the dscli command prompt with the parameters and variables shown as follows:

```
dscli>lssession -dev storage_image_ID -l LSS_ID  
dscli>lssession -dev IBM.2107-75FA120 -l 01
```

When you use the **-l** parameter, a detailed report displays a list of Global Mirror sessions for the specified logical subsystem (LSS) and information regarding 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 the Global Mirror session be paused and the parameters changed. Then you 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. Then start Global Mirror processing.

The tuning parameters of a Global Mirror session consist of the following:

- 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 observance of Global Mirror processing indicates that these parameters should be adjusted.
- After you have stopped Global Mirror processing as opposed to pausing Global Mirror processing. When you restart Global Mirror processing, the values for the tuning parameters revert to their DS CLI default values.

Perform the following steps to modify the tuning parameters of a Global Mirror session. 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 values declared 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  
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 writes to the target have been processed.

2. Issue the **showgmir** command to determine what the current 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, use: `dscli>showgmir -dev storage_image_ID -fullid LSS_ID`

For a performance metrics report, use: `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
```

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 3 tuning parameters. Enter the **resumegmir** command at the dscli command prompt with the following parameters and variables:

Note: The example command shows all 3 tuning parameters with new values. You must specify a value for all 3 tuning parameters even if only one value changed. The values for the two unchanged tuning parameters would be the DS CLI default values.

```
dscli>resumegmir -dev storage_image_ID -lss ID -cginterval seconds  
-coordinate milliseconds -drain seconds -session session_ID  
Master_Control_Path_LSS_ID:Subordinate_Control_Path_LSS_ID  
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 time writes are inhibited to the remote site, in seconds. The value indicates the maximum amount of time that writes are inhibited to the remote site before stopping the current consistency group. 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 which 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 handled differently.

Topology in this process refers to the list of storage servers which are subordinates. You establish remote mirror and copy paths between the Master and Subordinate

LSSs. Just one LSS per subordinate 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 go into the topology specification which defines the communication paths between Master and subordinate storage servers. To change these values you must stop the Global Mirror process.

Note: When you restart Global Mirror processing, the tuning parameters revert to their DS CLI default values.

Perform the following steps to modify the topology of a Global Mirror session. 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 **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 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.
- 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. The use of **Ctrl C** stops the DS CLI session. It does not stop the microcode 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 changes to the topology values 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 still need to issue the **mkgmir** command to resume Global Mirror processing when you have stopped the processing. You cannot issue the **resumegmir** command to start the processing. The **resumegmir** command is used in conjunction only with the **pausegmir** command.
- b. You must enter values for the tuning parameters (**-cginterval**, **-coordinate**, **-drain**) when you restart or start Global Mirror processing.

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.

The **lssession** command is used to display Global Mirror session information regarding the volumes in each session.

Perform the following step to view the associated properties of the Global Mirror session or information about Global Mirror session failures. 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.

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 **-l** parameter if you want to see a detailed report. The report provides you with information like the following:
 - State of the session status. For example, 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) 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.

The **showgmir** command to display a detailed report about the current Global Mirror operations and you can also request that the report include the detailed performance metrics.

Perform the following step to view the detailed properties and performance metrics for a Global Mirror session. 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.

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 parameters and variables shown as follows:

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 action 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

Perform the following steps to pause 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 **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, you should 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.

Perform the following steps to resume 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.

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>resumgmir -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 would look like the following:

```
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.

Perform the following step to start 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.

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 need to change the Global Mirror topology parameters, you need to 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. The steps in this task should only be taken if there is no alternative. For example a disaster has occurred and you need to 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. Of course this stops your DS CLI session. It is likely that this action might cause some transactions to remain partly completed and others completely 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 as these are related to the failover and failback tasks described in the *Recovering when a disaster strikes* scenario.

Perform the following step when you need 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 described in the *Recovering when a disaster strikes* scenario.

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.

Perform 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 similar to the following appears when you process the **rmgmir** command.

```
Are you sure you want to stop Session ID 01? [y/n]: Y
Global Mirror for Session ID 01 successfully stopped.
```

Setting up the Global Mirror Environment

This task lists the high-level steps that you must complete in order 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 169 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 171 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 172 for additional information.
4. Create the Global Mirror session. See “Creating the Global Mirror session” on page 173 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 is the first step in setting up your Global Mirror environment.

Create paths so that the logical subsystems (LSS's) 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.

Perform 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.

1. 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 -l
```

This is the entire command. there are no additional variables needed.

The **showsi** command does contain a variable and a command flag:

```
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** DS CLI command flag 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 **-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 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. For example, your command line would look like the following:

```
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 **-src~~l~~ss** 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 **-src~~l~~ss *source_LSS_ID*** value. For example: **-src~~l~~ss IBM.2107-75FA120/00**.
- d. The shortened version of the **-tgt~~l~~ss** 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 `-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 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 -src1ss 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 the fibre-channel paths between all Global Mirror source and target pairs and between the Master and subordinate storage units have been created.

The purpose of this step is to create a relationship between a source volume and a target volume.

Perform the following steps to create Global Copy pairs from the local storage units to the remote 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 values declared for the variables.

1. Issue the `mkpprc` command to create a Global Copy pairs 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 parameters and variables shown as follows:

```
: dscli>mkpprc -dev storage_imag_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. I/O write completion status is returned to the application when the updates are committed to the source.
- 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 .

The use of this parameter is positional. It must be placed after the **-type** parameter. For example:

```
dscli>mkpprc -type gcp IBM.2107-75FA120/0100:IBM.2107-75FA150/0100
```

A confirmation message is issued for each successful Global Copy volume association that is created.

2. Issue the **lspprc** command to check the status information for each remote mirror and copy volume relationship in the list. A report is displayed and you should see that each Global Copy relationship you have created has a status of copy pending. 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 -l 0100:0100
```

Use the **-l** parameter to provide a more detailed report about the Global Copy volume relationships.

It is recommended that 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 started. So you need to add the volumes to the session, establish the pairs, let the first pass complete, establish the FlashCopy relationships (BACKGROUND NOCOPY, change recording, inhibit target write), and then start the session.

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.

- 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.

The purpose of this step is to create a FlashCopy target for the Global Mirror pairs.

Perform 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 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 values declared for the variables.

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 as follows:

The **mkflash** command must contain the following parameters and variables as follows: 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 for the **-dev** *storage_image_ID* parameter. If the management console has an IP connection to the specified "remote async site" the command works. If the IP connection is not established you must use the **mkremoteflash** command with all the same parameters as displayed in the example.
- 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.

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 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 following parameters and variables:

```
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 used the **mkremoteflash** command, you must enter the **lsremoteflash** command to perform 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.

- 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 to associate volumes to a Global Mirror session.

Perform the following steps to create the Global Mirror session. 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 values declared for the variables.

1. Issue the **mksession** command to create the Global Mirror session. Enter the **mksession** command at the dscli command prompt using the parameters and variables shown as follows:

```
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 to 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 in order to remove the Global Mirror environment from your system.

The removal of a Global Mirror environment is prompted by circumstances similar to the following:

- You decide that you want to run with a totally 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 move on 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 175 for additional information.
3. Withdraw all FlashCopy relationships between the B and C volumes. See “Removing FlashCopy relationships” on page 176 for additional information.
4. Remove the Global Copy pair relationships. See “Removing the Global Copy pair relationship” on page 177 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 178 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.

Perform 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 the command requires. The second format provides the command with declared values for the variables.

1. Issue the **lsession** command to query the status of all volumes that are associated with the sessions of a specified logical subsystem. Enter the **lsession** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsession -dev storage_image_ID -l LSS_ID  
dscli>lsession -dev IBM.2107-75FA120 -l 01
```

When you use the **-l** parameter, a detailed report displays a list of Global Mirror sessions for the specified logical subsystem (LSS) and information regarding the volumes of each session in the list.

2. Analyze the report and determine which volumes 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 parameters and variables shown as follows:

```
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 closed. Ending a session does not remove the volume or path associations. Each of these must be removed in their own way.

Perform the following steps to close a Global Mirror session. 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 values declared for the variables.

1. Issue the **lssession** command to obtain a list of all sessions associated with the specified logical subsystem. Enter the **lssession** command at the dscli command prompt with the parameters and variables shown as follows:

```
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 causes a report with the following 3 pieces of information to display:

- LSSID
- Session number
- Volume numbers

2. Print the report or record the session numbers that need to be closed.
3. Issue the **rmsession** command to close the specified session. Enter the **rmsession** command at the dscli command prompt with the parameters and variables shown as follows:

```
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 you want the specified session stopped. A message similar to the following appears when you process the **rmsession** command.

```
Are you sure you want to stop Global Mirror Session ID 01? y/n Y
Global Mirror Session ID 01 successfully stopped
```

5. Repeat Step 3 for each session that needs to be closed.

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 invoking 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 perform a backup data before proceeding with this task.

Perform 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.

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 perform 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 similar to the following 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 initiating this task otherwise this task fails:

- 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.
- Remove FlashCopy relationships that existed 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 involved and the Global Copy relationships must be removed within each LSS. This requires that you repeat the remove command for each LSS until all relationships are removed.

Perform the following steps to remove Global Copy pair 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 values declared for the variables.

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 -l 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 parameters and variables shown as follows:

```
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 that requests that you confirm that you want the specified Global Copy pair removed. A message similar to the following appears for each Global Copy pair 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 Step 3 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 you use when you remove the Global Mirror environment from your system.

Ensure that the following tasks have been completed before you initiate this task, otherwise this task fails:

- Remove the volumes that are associated with each Global Mirror session you have closed.
- Close the Global Mirror sessions that are part of your Global Mirror environment.

- Remove FlashCopy relationships that existed with the volumes that were part of your Global Mirror environment.
- Remove the Global Copy relationships that existed with the volumes that were part of your Global Mirror environment.

You need to repeat the process described for issuing the remove command (**rmpprcpath**) for each data path (between LSSs on the source Storage Images on the primary site and the target LSSs on the target Storage Images on secondary site) and for each control path (between the Master Storage Image LSS and the Subordinate Storage Image LSSs).

Perform 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 the command requires. The second format provides the command with values declared 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 parameters and variables shown as follows:

```
dscli>lspprcpath -dev storage_image_ID source_LSS_ID
```

Example

```
dscli>lspprcpath -l -dev IBM.2107-75FA120 01
```

Note: The report displayed from this command provides the worldwide node name (WWNN) that is used with the **rmpprcpath** command.

2. Record the path information for use when issuing 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 parameters and variables shown as follows:

```
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 that requests that you confirm that you want the specified remote mirror and copy path removed. A message similar to the following appears for each remote mirror and copy path 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.
```

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 6. Recovering from a disaster using Global Mirror

This task lists the high-level steps that you must complete in order 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 perform the following high-level tasks using the Global Mirror function and the DS CLI commands:

1. End Global Mirror processing when a disaster occurs.
See *“Ending Global Mirror processing when a disaster occurs”* on page 182 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 183 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 184 for additional substeps.
4. Analyze the consistency group status. See *“Analyzing and validating the consistency group state”* on page 185 for additional substeps.
5. Use the **revertflash** command to correct FlashCopy relationships. See *“Using the revertflash command to correct FlashCopy relationships”* on page 187 for additional substeps.
6. Use the **commitflash** command to correct FlashCopy relationships. See *“Using the commitflash command to correct FlashCopy relationships”* on page 188 for additional substeps.

7. Initiate the fast reverse restore process. See *“Using fast reverse restore processing to create consistency” on page 189* for additional substeps.
8. Wait for the background copy to complete. See *“Waiting for the background copy to complete” on page 190* 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 191* for additional substeps.
10. Prepare to reinstate production at the local site. See *“Preparing to reinstate production at the local site” on page 192* for additional substeps.
11. Resynchronize the volumes. See *“Resynchronizing the volumes” on page 193* for additional substeps.
12. Query for first pass and drain time out-of-synch zero value and quiesce your system. See *“Querying, quiescing, and querying” on page 194* 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 195* for additional substeps.
14. Perform Global Copy failover processing to the A volumes. See *“Performing Global Copy failover processing to the A volumes” on page 197* for additional substeps.
15. Perform Global Copy failback processing for the A volumes. See *“Performing Global Copy failback processing for the A volumes” on page 197* for additional substeps.
16. Resume Global Mirror processing at site A. See *“Resuming Global Mirror processing at site A” on page 198* 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, when possible.

Depending on the state of the local Global Mirror storage server, you might have an opportunity to end Global Mirror processing before initiating the rest of the recovery steps. Generally a disaster affects your local server and your choices for ending Global Mirror processing in one of the following ways:

- You experience a rolling disaster and your Global Mirror processing is not being done through a DS CLI script. This circumstance allows you time to initiate some planning by issuing a pause command, followed by a query command, and then issuing the 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.

- You experience 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 will have to 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 local server is affected immediately and you have no time to end Global Mirror processing or issue a status query.

When possible, you want to end Global Mirror processing. Perform 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 was ended as a result of the disaster. This is the second step in the Global Mirror failover recovery process when you have time because of dealing with a rolling disaster and 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 initiate this task.

You must determine where Global Mirror processing was before processing was ended. Some transactions might be half completed while others were not yet started. Querying the status of your transactions provides a basis for planning what must be done next. The situation you face is similar to the following:

- The formation of a consistency group was in progress and the FlashCopy state between the B volumes and the C volumes in the remote storage server are not the same at the for all relationships when Global Mirror processing was ended.
- Some FlashCopy pairs might have completed the FlashCopy phase of Global Mirror processing to form a new consistency group and already committed the changes.
- Some FlashCopy pairs might not have completed yet 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 what might still go on for a brief period at the remote site.

Perform the following steps to obtain an initial idea of the status of your transactions.

1. 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 still operational. 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 application.
3. Log in to a DS CLI session.
4. Issue the **lspprc** 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 -l
0100-0103:0100-0103 0200:0200 0300-0303:0300-0303
```

Note: The report displays your Global Copy pairs with a suspended state.

After you have 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 initiate 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 effect is that the Global Copy state of the B volumes changes from secondary to primary and suspended. This is the next step after obtaining a status of your Global Mirror transactions in the Global Mirror failover recovery process.

The failover task cannot be done unless you have completed the following:

- Ended Global Mirror processing.
- Obtained the report or reports that help you determine the status of your Global Mirror transactions before the disaster occurred.

You must issue a recovery failover request on the Global Copy volumes pair to force a stop of the A Volume to B Volume extended distance relationship and create a B Volume to A Volume Global Copy relationship. 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.

Perform the following step to fail over the B volumes to A volumes. 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.

Issue the **failoverpprc** command to change the Global Copy state on the B volumes from secondary, target pending to primary, suspended. 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  
0100:0100 0101:0101 0102:0102 0103:0103
```

Note: Unlike most other commands, the values for the *storage_image_ID* and *SourceVolumeID:TargetVolumeID* are reversed in the failoverpprc command. That is, the command is issued directly to the remote server. So, the values for the remote server are first and the values for the original primary server are second. A role reversal is affected so that the remote becomes the primary server. This results in the following:

- The source volumes show as suspended host.
- The target volumes show as suspended target and they are accessible for mounting.

A confirmation message like the following is generated for each Global Copy pair that has been changed and moved to a state of suspended.

```
PPRC pair IBM.2107-75FA150/0100:IBM.2107-75FA120/0100 successfully
suspended.
```

Note: All B volumes must successfully process the **failoverpprc** command before you can move on to the next step.

After all the B volumes have been processed, you are ready for the next step which involves analyzing the FlashCopy relationships and initiating 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 so as 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.

Perform 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.

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 at this point 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 new 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 very 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 exist:
 - 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.

Perform 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.

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 performed 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 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 new 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 were involved in the very 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 of before the disaster occurred.

Perform 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.

Issue the `commitflash` command to correct the FlashCopy relationships and commit them 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 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 **commitflash** 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 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.

Perform 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.

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. The source volume ID is the value that is specified for the C volume. The data from this volume is copied to the target 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 go on 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 background 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.

Perform 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.

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.

Perform 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.

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.

Perform 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.

Note: 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).

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 -src1ss source_LSS_ID -tgt1ss target_LSS_ID  
source_port_ID:target_port_ID
```

Example

```
dscli>mkpprcpath -dev IBM.2107-75FA150 -remotedev  
IBM.2107-75FA120 -remotewwnn 12341234000A000A  
-src1ss IBM.2107-75FA150/00 -tgt1ss 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 will discover which tracks were modified at any site and send both the tracks that were changed on Site A and the tracks that were marked at Site B from Site A to Site B.

Perform 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.

Note: 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.

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 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 perform 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.

Perform 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.

1. Issue the **lspprc** 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 period of 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
```

Notes:

- a. The **-dev** parameter specifies the ID of your source storage unit. At this point, your source is the Site B storage unit.
 - b. The **-remotedev** parameter specifies the ID of the target storage unit. At this point, 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 550 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 period of 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 completed in sequence in order for this task to succeed.

This task is similar to when you initially created your remote mirror and copy paths as part of 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.

Perform 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.

1. 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 -l
```

This is the entire command. No additional variables are needed.

The **showsi** command does contain a variable and a command flag:

```
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 **shows** command if you want to display the detailed properties of a specific storage unit.
 - c. Use the **-fullid** DS CLI command flag with the **shows** 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 **-src~~l~~ss** 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 **-src~~l~~ss *source_LSS_ID*** value. For example: **-src~~l~~ss IBM.2107-75FA120/00**.
- d. The shortened version of the **-tgt~~l~~ss** 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 **-tgt~~l~~ss *target_LSS_ID*** value. For example: **-tgt~~l~~ss 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 **-tgt~~l~~ss** parameter and value. For example:

```
dscli>mkpprcpath -srclss 00 -tgtlss 01
IBM.2107-75FA120/I1A10:IBM.2107-75FA150/I2A20
```

Performing Global Copy failover processing to the A volumes

Complete this task to perform 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 failback recovery 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.

Perform 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.

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 gcpSourceVolumeID: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 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 move on to the next step.

Performing Global Copy failback processing for the A volumes

Complete this task to perform 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 failback recovery process.

The failover processing that is described in the fifth step of the failback recovery 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 in order to resume mirroring.

Perform 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.

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.

Perform 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.

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 7. Recovery scenarios for planned and unplanned outages using Metro/Global Mirror

This section describes a Metro/Global Mirror environment that uses failover and failback operations to switch applications to an alternate site during planned outages such as maintenance updates or unplanned outages such as disasters:

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
- Failover and restore operations to the intermediate (B) site
- Failover and restore operations to the remote (C) site

Setting up a Metro/Global Mirror environment

Use this process to set up your system environment 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 will be 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.

The following steps are in the recommended sequence. You can perform 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.

Perform 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 the I/O ports are 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 154 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 -src1ss 61 -tgt1ss 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` option) 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 151 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 -src1ss 62 -tgt1ss 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 139 for more information.

d. Use the `lspprcpath` command to view the newly created paths. See “Displaying the status of established paths” on page 138 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` option 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 146 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 142 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: It is recommended that initial FlashCopy operations be delayed 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 172 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 173 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 160 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 167 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 165 for more information.

Enter the **showgmir** command at the dscli command prompt with the following parameters and variables:

```
dsccli> 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	Global Mirror Interval	CG Drain Time	Current Time
IBM.2107-75ALA2P/07	1	0x31	Running	Not Fatal	0	50	30	02/19/2006 10:34:51 MST

CG Time	CG Failed Attempt	Successful CG Percentage	Flash-Copy Sequence Number	Master ID	Subordinate Count	Master SSID	Subordinate SSID
02/19/2006 10:34:51 MST	0	50	0x4357D-504	IBM.2107-75ALA2P	0	-	-

Failover and restore operations to the intermediate site during a planned outage

Use this process to perform failover and restore operations 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 procedure 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 perform 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 perform 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 will help coordinate the A volumes and B

volumes consistency and allow consistent data to be copied to the remote site. You can use either one 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 207.
- Freeze write activity to the Metro Mirror primary volumes by performing the following steps:

- 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:

- 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 Metro Mirror volumes.
 - The established paths between the logical subsystem (LSS) pairs are disabled.
- If desired, 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 flag **-fmt default**. This command flag 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 **-l** parameter.

See “Viewing information about Metro Mirror relationships” on page 153 for more information.

ID	State	Reason	Type	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

- 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 other words, 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 “Performing a failover recovery operation” on page 152 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 165 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 145 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 139 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 “Performing a failback recovery operation” on page 152 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 153 for more information.

Enter the **lsprrc** command at the dscli command prompt with the following parameters and variables:

```
dscli> lsprrc -l 1200-125f:1a00-1a5f
```

The following represents an example of the output:

ID	State	Reason	Type	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 Suspend	Source LSS	Timeout (secs)	Crit Mode	First Pass Status	Incremental Resync	Tgt Write
Invalid	-	10	Unknown	Diabled	Invalid	Enabled	Enabled
Invalid	-	10	Unknown	Diabled	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 “Performing a failover recovery operation” on page 152 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:

- The command example uses the command flag **-fmt default**. This command flag specifies that the output be set to a space-separated plain text table.
- The following table format is presented for clarity. The actual report is not displayed in this format.
- The report example represents the information that is reported on when you do not specify the **-I** parameter.

ID	State	Reason	Type	Source-LSS	Timeout (secs)	Critical Mode	First Pass Status
0700:1200	Suspended	Host Source	Metro Mirror	1A	unknown	Disabled	Invalid
0701:1201	Suspended	Host Source	Metro Mirror	1A	unknown	Disabled	Invalid
0702:1202	Suspended	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 -src1ss 07 -tgt1ss 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 195 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 will be immediately started 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 will not be fully protected by Global Mirror processing until the Global Mirror operation is restored in step 16 on page 211.

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.
```

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 recreated.

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 -1
-fmt default 1200-125f:0700-075f
```

The following represents an example of the output:

ID	State	Reason	Type	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 Suspended	Source LSS	Timeout (secs)	Crit Mode	First Pass Status	Incremental 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 166 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 165 for more information.

ID	Total Failed CG Count	Total Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milliseconds)	Interval Time (seconds)
IBM.2107-130165X/07	0	55	100	0	02/20/2006 11:38:25 MST	50	0

Max 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 perform failover and restore operations 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, perform 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 actions processing occurs:

- I/O 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 in Metro Mirror relationships.
- The established paths between the LSS pairs are disabled.

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 “Performing a failover recovery operation” on page 152 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 165 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 145 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 -src1ss 61 -tgt1ss 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 139 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 “Performing a failback recovery operation” on page 152 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 153 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 “Performing a failover recovery operation” on page 152 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 195 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 will be immediately started 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 will not be fully protected by Global Mirror processing until the Global Mirror operation is restored in step 15 on page 216.

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 “Performing a failback recovery operation” on page 152 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 recreated. 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 146 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 166 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 Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milliseconds)	Interval Time (seconds)
IBM.2107-130165X/07	0	55	100	0	10/20/2005 11:38:25 MST	50	0

Max 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 165 for more information.

Failover and restore operations at the remote site during a planned outage

Use this process to perform failover and restore operations 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 219.

If you freeze write activity to the Metro Mirror primary volumes, perform 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 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:

- 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 desired, 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 flag **-fmt default**. This command flag 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 **-l** parameter.

ID	State	Reason	Type	Source-LSS	Timeout (secs)	Critical Mode	First Pass Status
0700:1200	Suspended	Freeze	Metro Mirror	07	unknown	Disabled	Invalid
0701:1201	Suspended	Freeze	Metro Mirror	07	unknown	Disabled	Invalid
0702:1202	Suspended	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 represents an example of the output:

```
CMUC00198I unfreezepprc: Remote Mirror and Copy pair 07:12
successfully thawed.
```

- 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
```

The following represents an example of the output:

ID	Total Failed CG Count	Total Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milliseconds)	Interval Time (seconds)
IBM.2107-75ALA2P/10	23	139	85	0	02/20/2006 11:33:56 MST	50	0

Max 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.2107-75ALA2P	0x12	Error	Session or Session Members not in Correct State	Global Mirror Run in Progress	IBM.2107-75ALA2P	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.2107-75ALA2P	Not Available	Error	Max Drain Time Exceeded	Drain in Progress

- 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 represents an example of the output:

```
CMUC00165I rmgmir: Global Mirror for session 31 successfully stopped.
```

See “Ending Global Mirror processing (script mode)” on page 167 or “Ending Global Mirror processing (no script)” on page 168 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 **rmgmir** 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 Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milliseconds)	Interval Time (seconds)
IBM.2107-75ALA 2P/10	414	148054	99	0	03/20/2006 15:33:56 MST	50	0

Max 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 The following represents an example of the output: (a shortened version) is displayed:

```
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 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.
```

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
```

The following represents an example of the output:

ID	State	Reason	Type	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 Suspend	Source LSS	Timeout (secs)	Crit Mode	First Pass Status	Incremental 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
-remotewwn 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 139 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 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 “Performing a failback recovery operation” on page 152 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 -l -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> rmprrc -quiet -dev IBM.2107-1831760 -remotedev IBM.2107-75ALA2P 0700-075f:1200-125f
```

The following represents an example of the output:

```
CMUC00155I rmprrc: Remote Mirror and Copy volume pair 0700:1200 relationship successfully withdrawn.
CMUC00155I rmprrc: Remote Mirror and Copy volume pair 0701:1201 relationship successfully withdrawn.
CMUC00155I rmprrc: 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 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.
```

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 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 139 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 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 “Performing a failback recovery operation” on page 152 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 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 146 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 represents an example of the output:

```
CMUC00137I mkflash: FlashCopy pair 1300:1900 successfully created.  
CMUC00137I mkflash: FlashCopy pair 1301:1901 successfully created.  
CMUC00137I mkflash: FlashCopy pair 1302:1902 successfully created.
```

See “Creating FlashCopy relationships (Global Mirror setup)” on page 172 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 represents an example of the output:

```
CMUC00164I resumegmir: Global Mirror for session 31 successfully resumed.
```

See “Resuming Global Mirror processing” on page 166 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 perform failover and restore operations 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 site (Site A) 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 Site C. This process is known as a failover recovery. The Global Copy relationship between volumes at the intermediate and remote site are still operational. Global Mirror continues to operate between these two sites.

Note: The steps to perform a failover operation to the remote site (C) can be done using a *single* GDPS automation script running in a GPDS Global Mirror remote controlling system. GDPS provides a recovery check function that determines the state of the volumes *before* necessary actions to perform a recovery function is done. This process alerts the user to fix the required problem before performing the actual recovery is done, which can reduce the time required to perform the recovery.

Perform 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, perform 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.) 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:

- 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. **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 Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milliseconds)	Interval Time (seconds)
IBM.2107-75ALA2P/10	23	139	85	0	02/20/2006 11:33:56 MST	50	0

Max 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.2107-75ALA2P	0x12	Error	Session or Session Members not in Correct State	Global Mirror Run in Progress	IBM.2107-75ALA2P	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.2107-75ALA2P	Not Available	Error	Max Drain Time Exceeded	Drain in Progress

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 167 or “Ending Global Mirror processing (no script)” on page 168 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 165 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 Interval	Global Mirror Interval	CG Drain Time	Current Time
IBM.2107-75ALA2P/10	-	-	-	-	-	-	-	-

CG Time	CG Failed Attempt	Successful CG Percentage	Flash-Copy Sequence Number	Master ID	Subordinate Count	Master SSID	Subordinate SSID
-	-	-	-	-	-	-	-

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 rmprrc: Remote Mirror and Copy volume pair 1200:0700 relationship
successfully withdrawn.
CMUC00155I rmprrc: Remote Mirror and Copy volume pair 1201:0701 relationship
successfully withdrawn.
CMUC00155I rmprrc: Remote Mirror and Copy volume pair 1202:0702 relationship
successfully withdrawn.
```

See “Deleting a Metro Mirror relationship” on page 147 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 “Performing a failover recovery operation” on page 152 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 146 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 172 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 173 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 167 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
```

The following represents an example of the output:

ID	Master Count	Master Session ID	Copy State	Fatal Reason	CG Inter-val	Global Mirror Inter-val	CG Drain Time	Current Time
IBM.2107-75ALA2P/07	1	0x01	Running	Not Fatal	0	50	30	02/20/2006 11:37:40 MST

CG Time	CG Failed Attempt	Succes-sful CG Percen-tage	Flash-Copy Sequ-ence Number	Master ID	Subor-dinate Count	Master SSID	Subor-dinate SSID
02/20/2006 11:37:40 MST	-	100	0x4357E-3F4	IBM.2107-75ALA2P/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	Global Mirror Inter- val	CG Drain Time	Current Time
IBM. 2107- 75ALA 2P/07	1	0x01	Running	Not Fatal	0	50	30	02/20/ 2006 11:37:40 MST

CG Time	CG Failed Attempt	Succes- ful CG Percen- tage	Flash- Copy Sequ- ence Number	Master ID	Subor- dinate Count	Master SSID	Subor- dinate SSID
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 will be 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 “Performing a failback recovery operation” on page 152 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 -1
-fmt default 1200-125f
```

The following represents an example of the output:

ID	State	Reason	Type	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 Suspend	Source LSS	Timeout (secs)	Crit Mode	First Pass Status	Incremental Resync	Tgt Write
Invalid	-	12	Unknown	Diabled	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 167 or “Ending Global Mirror processing (no script)” on page 168 for more information.

16. **Verify that at least two consistency groups have formed.** Assuming that consistency groups formed successfully, the A, B, C, and E volumes will 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 165 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 Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milliseconds)	Interval Time (seconds)
IBM.2107-75ALA2P/10	0	55	100	0	02/20/2006 11:38:25 MST	50	0

Max 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	-	-

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:

```
CMUC00165I rmgmir: Global Mirror for session 1 successfully stopped
```

See “Ending Global Mirror processing (script mode)” on page 167 or “Ending Global Mirror processing (no script)” on page 168 for more information.

18. **Verify that the Global Mirror session for which the C, B, and E volumes are included, has stopped.** Enter the **showgmir** command at the dscli command prompt with the following parameters and variables.

```
showgmir 07
```

The following represents an example of the output:

ID	Total Failed CG Count	Total Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milliseconds)	Interval Time (seconds)
IBM.2107-75ALA 2P/10	23	139	85	0	02/20/2006 11:33:56 MST	50	0

Max 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 165 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 174 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, which is described in step 24 on page 234. 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 147 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 195 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 146 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) at 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 172 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 176 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 166 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 Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milliseconds)	Interval Time (seconds)
IBM.2107-75ALA2P/10	1	39	97	0	02/20/2006 11:33:56 MST	50	0

Max 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.2107-75ALA2P	0x12	Error	Session or Session Members not in Correct State	Global Mirror Run in Progress	IBM.2107-75ALA2P	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 165 for more information.

Using forced failover and failback during a planned Metro/Global Mirror outage

Use this process to perform failover operations 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. For additional information, refer to Creating a new Global Mirror session in the Related tasks section below.

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 task 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.

Perform 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, perform 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:

```
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.

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.
```

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.

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. (In other words, 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	Type	Source-LSS	Time-out (secs)	Critical Mode	First Pass Status
IBM.2107-183176O /2101: IBM.2107-130165X /2101	Copy Pending	-	Global Copy	IBM.2107-130165X /20	300	Disabled	True
IBM.2107-183176O /2100: IBM.2107-130165X /2100	Copy Pending	-	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 IBM2107-183176O -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 Global Mirror has terminated.

Enter the **showgmir** command at the dscli command prompt with the following parameters and variables:

```
dscli> showgmir -dev storage_image_ID 10
```

Example

```
dscli> showgmir -dev IBM.2107-75ALA2P
```

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 -l  
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 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 -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 -remotedev storage_image_ID -dev storage_image_ID
-type mmir -mode nocp -incrementalresync enable SourceVolumeID:TargetVolumeID
```

Example

```
dscli>mkpprc -remotedev IBM.2107-1301261 -dev IBM.2107-130165X -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.2107-75ALA2P
-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 **lsprrc** command if you want to monitor this activity and determine when the first pass status changes to "True."

Enter the **lsprrc** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsprrc -dev storage_image_ID -remotedev storage_image_ID
-l SourceVolumeID:TargetVolumeID
```

Example

```
dscli>lsprrc -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 **lsprrc** command to query the status of the C to B volume pairs in Global Copy relationships.

Enter the **lsprrc** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsprrc -dev storage_image_ID -remotedev storage_image_ID
-l SourceVolumeID:TargetVolumeID
```

Example

```
dscli> lsprrc -dev IBM.2107-75ALA2P -remotedev IBM.2107-1831860 -l
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 -dev storage_image_ID -remotedev storage_image_ID
-type gcp 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 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 -1  
-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 Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milliseconds)	Interval Time (seconds)
IBM.2107-130165X/07	0	55	100	0	10/20/2005 11:38:25 MST	50	0

Max 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	-	-

Using forced failover and failback during an unplanned Metro/Global Mirror outage

Use this process to perform failover operations 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. For additional information, refer to Creating a new Global Mirror session in the Related tasks section below.

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.

Perform 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, perform 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:

- I/O 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 in Metro Mirror relationships.
- The established paths between the LSS pairs are deleted.

- 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 using the -force parameter from the C volumes to the A volumes.** No validation is done at site C to determine that site C is a secondary of site A.

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 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.

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. (In other words, 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 **lsprrc** command:

ID	State	Reason	Type	Source-LSS	Time-out (secs)	Critical Mode	First Pass Status
IBM.2107-183176O /2101: IBM.2107-130165X /2101	Copy Pending	-	Global Copy	IBM.2107-130165X /20	300	Disabled	True
IBM.2107-183176O /2100: IBM.2107-130165X /2100	Copy Pending	-	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 Global Mirror has terminated.

Enter the **showgmir** command at the dscli command prompt with the following parameters and variables:

```
dscli> showgmir -dev storage_image_ID 10
```

Example

```
dscli> showgmir -dev IBM.2107-75ALA2P
```

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 -l
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 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.2107-75ALA2P -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.

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 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 -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 -remotedev storage_image_ID -dev storage_image_ID  
-type mmir -mode nocp -incrementalresync enable SourceVolumeID:TargetVolumeID
```

Example

```
dscli>mkpprc -remotedev IBM.2107-1301261 -dev IBM.2107-130165X -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 using the -force parameter from the C volumes to the A volumes.**

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.
```

26. **Wait for the first pass copy to complete from site C to site B.** Issue the **lsprrc** command if you want to monitor this activity and determine when the first pass status changes to "True."

Enter the **lsprrc** command at the dscli command prompt with the following parameters and variables:

```
dscli>lsprrc -dev storage_image_ID -remotedev storage_image_ID
-l SourceVolumeID:TargetVolumeID
```

Example

```
dscli>lsprrc -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 -l  
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 with cascading 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 -dev storage_image_ID -remotedev storage_image_ID
-type gcp 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 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 Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milliseconds)	Interval Time (seconds)
IBM.2107-130165X/07	0	55	100	0	10/20/2005 11:38:25 MST	50	0

Max 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 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 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 perform 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, it is recommended that you 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 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.
```

See “Creating a Metro Mirror relationship” on page 142 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 represents an example of 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 145 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 this is the case, you must verify the consistency group formation. See “Querying Global Mirror processing” on page 165 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 represents an example of 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 “Performing a failover recovery operation” on page 152 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 represents an example of the output:

ID	State	Reason	Type	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 Cascade	Date Suspended	Source LSS	Time-out (secs)	Crit Mode	First Pass Status	Incremental 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 represents an example of the output:

```
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 175 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 126 for more information.

The following represents an example of the output:

ID	SrcLSS	Sequence 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	Background 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

- 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:

```
dscli> mkpprc -dev IBM.2107-130165X -remotedev IBM.2107-1831760 -type gcp
-incrementalresync recover 2100-2107:2100-2107
```

See “Creating a Global Copy relationship” on page 146 for more information.

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.
```

Notes:

- 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.
 - 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.
- 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.

From the local site, enter the `lspprc` command at the `dscli` command prompt with the following parameters and variables:

```
dscli> lspprc -fullid 2100-2107
```

The following represents an example of the output:

ID	State	Reason	Type	Source-LSS	Time-out (secs)	Critical Mode	First Pass Status
IBM.2107-130165X /2101: IBM.2107-1831760 /2101	Copy Pending	-	Global Copy	IBM.2107-1831760 /20	300	Disabled	True
IBM.2107-130165X /2100: IBM.2107-1831760 /2100	Copy Pending	-	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 represents an example of the output:

```
CMUC00162I mkgmir: Global Mirror for session 31 successfully started.
```

See “Starting Global Mirror processing” on page 167 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. Perform 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 represents an example of 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 147 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 -fullid 2100-2107
```

The following represents an example of the output: See “Pausing a Metro Mirror relationship” on page 145 for more information.

ID	State	Reason	Type	Source-LSS	Time-out (secs)	Critical Mode	First Pass Status
IBM.2107-1831760 /2100: IBM.2107-75ALA2P /2100	Suspended Host Source	-	Global Copy	IBM.2107-75ALA2P /21	unknown	Disabled	True
IBM.2107-1831760 /2101: IBM.2107-75ALA2P /2101	Copy Pending	-	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 represents an example of the output:

```
CMUC00165I rmgmir: Global Mirror for session 31 successfully stopped.
See “Ending Global Mirror processing (script mode)” on page 167 or
“Ending Global Mirror processing (no script)” on page 168 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 represents an example of the output:

```
CMUC00165I rmgmir: Global Mirror for session 31 successfully
stopped.
```

See “Ending Global Mirror processing (script mode)” on page 167 or “Ending Global Mirror processing (no script)” on page 168 for more information.

11. **From the remote site, perform 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 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.
```

See “Performing a failback recovery operation” on page 152 for more information.

- 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 165 for more information. The following represents an example of the output:

ID	State	Reason	Type	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

Tgt Cascade	Date Suspend	Source LSS	Time-out (secs)	Crit Mode	First Pass Status	Incremental 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 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.
```

See "Creating a Metro Mirror relationship" on page 142 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 represents an example of the output:

```
CMUC00165I rmgmir: Global Mirror for session 31 successfully stopped.
```

See "Ending Global Mirror processing (script mode)" on page 167 or "Ending Global Mirror processing (no script)" on page 168 for more information.

15. **Allow the resynchronization of the C to B volumes to complete by performing 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 represents an example of 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 145 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 -l -fullid  
-fmt default 2100-2107
```

See "Querying Global Mirror processing" on page 165 for more information.

The following represents an example of the output:

ID	State	Reason	Type	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

Tgt Cascade	Date Suspend	Source LSS	Time-out (secs)	Crit Mode	First Pass Status	Incremental 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 represents an example of 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 177 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 represents an example of 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 145 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 gpc
-cascade 2100-2107:2100-2107
```

The following represents an example of 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 “Performing a failover recovery operation” on page 152 for more information.

18. **At the intermediate site, perform 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 gpc
-cascade 2100-2107:2100-2107
```

The following represents an example of 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 “Performing a failback recovery operation” on page 152 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 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.
```

See “Creating a Metro Mirror relationship” on page 142 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 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.
```

See “Creating a Metro Mirror relationship” on page 142 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 153 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 represents an example of the output:

```
CMUC00162I mkgmir: Global Mirror for session 31 successfully started.
```

See “Starting Global Mirror processing” on page 167 for more information.

Your original configuration is restored.

Chapter 8. CLI commands

This section describes the command-line interface (CLI) commands that you can use to perform configuration and storage management tasks.

About CLI commands

This is 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 perform. 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 perform the command task in a specific way. For example, the *-v* flag tells the command-line interface to display the command results in verbose mode. Some flags may 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 may 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 will have 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 if you specify *-v* only, then the flag parameter is assumed to be on.

The **command parameter** provides basic information that is necessary to perform 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 **lsuser**, do not require a command parameter, because a default value of *all* is always assumed. For some commands, like **lsarraysite**, the command parameter is optional. If no value is provided, then a default value of *all* is assumed. If a value is provided, then the command-line interface lists information only about the array site or sites provided in the command parameter string.

In the following example, **lsrank** is the command name. **-dev** and **-l** are command parameters. *IBM.2107-75FA120* is the sub parameter for the **-dev** parameter, and *R1*, *R2*, and *R3* are a list of command parameters.

```
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



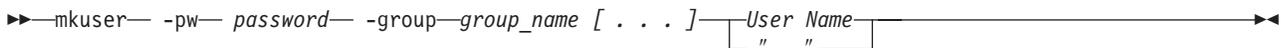
Begins on the left with double arrowheads (>>) and ends on the right with two arrowheads facing each other (><). If a diagram is longer than one line, each line to be continued ends with a single arrowhead (>) and the next line begins with a single arrowhead. Read the diagrams from left-to-right, top-to-bottom,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

Keyword



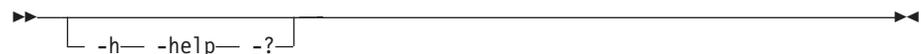
Represents the name of a command, flag, parameter, or argument. A keyword is not in italics. Spell a keyword exactly as it is shown in the syntax diagram.

Required keywords



Indicate 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



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

Variable



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

Special characters

- (minus) or / (slash) sign

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

[] square brackets

Optional values are enclosed in square brackets.

{ } braces

Required or expected values are enclosed in braces.

| vertical bar

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

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

... ellipsis

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

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

List of commands

This is a complete list of the command-line interface commands, alphabetized by command name.

Command	Type	Description
applykey	application key and version	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. The file that contains the LMC keys must be downloaded from an IBM Web site.
chckdvol	storage configuration	The chckdvol command changes the name of a count key data (CKD) base volume.
chextpool	storage configuration	The chextpool command modifies an extent pool.
chfbvol	storage configuration	The chfbvol command changes the name or data type of a fixed block volume.
chhostconnect	I/O port and host connect configuration	The chhostconnect command modifies a SCSI host port configuration.
chlcu	storage configuration	The chlcu command modifies a logical control unit.

Command	Type	Description
chlss	storage configuration	The chlss command modifies a logical subsystem.
chpass	user account and security	The chpass command changes the password expiration time and the number of login attempts for a storage complex.
chrank	storage configuration	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.
chsession	Copy Services	The chsession command allows you to modify a Global Mirror session.
chsi	storage image configuration	The chsi command modifies a storage image. You can use it to set characteristics such as online or offline state, name, and description.
chsp	storage complex configuration	The chsp command modifies a storage complex for items such as notification of the Simple Network Management Protocol (SNMP) traps and e-mail problem notification lists in a storage complex.
chsu	storage unit configuration	The chsu command modifies a storage unit.
chuser	user account and security	The chuser command is used to modify and lock or unlock a DS CLI or a DS Storage Manager user account. A CLI user with administrative authority uses this command to update a user account password, modify user group authority, or to lock or unlock a user account. Users that 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.
chvolgrp	storage configuration	The chvolgrp command modifies a volume group name and volume members.
clearvol	storage configuration	The clearvol command clears Copy Services relationships for a base logical volume.
commitflash	Copy Services	The commitflash command is used as part of the recovery from a disaster scenario to complete a partially formed Global Mirror consistency group.

Command	Type	Description
commitremoteflash	Copy Services	The commitremoteflash command sends data to a target volume to form a consistency between the remote source and target FlashCopy pair.
dscli	framework	The dscli command starts DS CLI. Use this command to run DS CLI commands in the interactive, single-shot, or script mode.
failbackpprc	Copy Services	The failbackpprc command copies the required data from the source volume to the target volume in order 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.
failoverpprc	Copy Services	The failoverpprc command is used only with disaster recovery processing. This command is used in the disaster recovery processes associated with sites using 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.
freezepprc	Copy Services	The freezepprc command creates a new remote mirror and copy consistency group. It places the source logical subsystem (LSS) in the <i>long busy</i> 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 <i>queue full</i> condition for the primary volume. This causes the host to queue writes to the primary volume until the <i>queue full</i> condition is reset. During the <i>queue full</i> condition, the primary volume reports <i>long busy</i> status.
lsaddressgrp	storage configuration	The lsaddressgrp command displays a list of address groups for a storage image and the status information for each address group in the list.
lsarray	storage configuration	The lsarray command displays a list of arrays in a storage image and status information for each array in the list.

Command	Type	Description
lsarraysite	storage configuration	The lsarraysite command displays a list of array sites and status information for each array site in the list.
lsavailpprcport	Copy Services	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 supports only fibre-channel ports. The Enterprise Storage Server (2105 machine type) supports ESCON ports.
lsckdvol	storage configuration	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.
lsda	physical resource information	The lsda command displays a list of device adapters (DA) for each storage image. You can use this command to look at the status of each device adapter in the list.
lsddm	physical resource information	The lsddm command displays a list of disk drive modules (DDMs) and status information for each DDM in the list.
lshba	physical resource information	The lshba command displays a list of storage image host adapters and status information for each host adapter in the list.
lsxtpool	storage configuration	The lsxtpool command displays a list of extent pools in a storage unit and status information on each extent pool in the list.
lsfbvol	storage configuration	The lsfbvol command displays a list of fixed block volumes in a storage image and status information for each volume in the list.
lsflash	Copy Services	The lsflash command displays a list of FlashCopy relationships and status information for each FlashCopy relationship in the list.
lsframe	physical enclosure information	The lsframe command displays a list of frame enclosures for a storage image.

Command	Type	Description
lshostconnect	I/O port and host connect configuration	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 new host connection using the mkhostconnect command.
lshosttype	I/O port and host connect configuration	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.
lshostvol	I/O port and host connect configuration	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.)
lsioport	I/O port and host connect configuration	The lsioport command 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.
lskey	application key and version	The lskey command displays the type of licensed machine code (LMC) activation keys that are installed and available for use by the storage unit.
lslcu	storage configuration	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.
lslss	storage configuration	The lslss command displays a list of logical subsystems (LSSs) for a storage image and status information for each logical subsystem in the list.
lsnetworkport	Network ports	The lsnetworkport command (DS8000 only) displays a report that lists all network ports that are installed in the specified storage unit.

Command	Type	Description
lsportprof	I/O port and host connect configuration	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.
lspprc	Copy Services	The lspprc command displays a list of remote mirror and copy (formerly PPRC) volume relationships for a storage image, and status information for each remote mirror and copy volume relationship in the list.
lspprcpath	Copy Services	The lspprcpath command displays a list of existing Remote Mirror and Copy (formerly PPRC) path definitions.
lsrank	storage configuration	The lsrank command displays a list of defined ranks in a storage image and status information for each rank.
lsremoteflash	Copy Services	The lsremoteflash command displays a list of FlashCopy relationships and status information for each FlashCopy relationship in the list.
lsserver	storage image configuration	The lsserver command displays all servers in a storage complex or a list of specified servers and it also displays the status information for each server in the list.
lssession	Copy Services	The lssession command displays a list of Global Mirror sessions for a logical subsystem (LSS) and information regarding the volumes of each session in the list.
lssi	storage image configuration	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.
lststgencl	physical enclosure information	The lststgencl command displays a list of storage enclosures and status information for each enclosure in the list.

Command	Type	Description
lssu	storage unit configuration	The lssu 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.
luser	user account and security	The luser command returns a list of storage image user account names and access authority levels.
lsvolgrp	storage configuration	The lsvolgrp command displays a list of volume groups in a storage image and status information for each volume group in the list.
managehostconnect	I/O port and host connect configuration	The managehostconnect command modifies the volume group assignment for a SCSI host port.
managepwfile	user account and security	The managepwfile command creates a password file for an existing ESS or DS user account. This command processes the password requirements for 2105, 2107, and 1750 systems.
mkaliasvol	storage configuration	The mkaliasvol command creates zSeries CKD alias volumes (generally referred to as parallel access volumes or PAVs) in a storage image.
mkarray	storage configuration	The mkarray command creates one array per command.
mkckdvol	storage configuration	The mkckdvol command creates zSeries count key data (CKD) base or CKD alias volumes in a storage image.
mkesconpprcpath	Copy Services	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 System Storage Enterprise Storage Servers (2105, Model 800 and Model 750).
mkextpool	storage configuration	The mkextpool command creates a fixed block or count key data (CKD) storage type extent pool.
mkfbvol	storage configuration	The mkfbvol command creates open systems fixed block (FB) volumes in a storage image.
mkflash	Copy Services	The mkflash command initiates a point-in-time copy from source volumes to target volumes.

Command	Type	Description
mkgmir	Copy Services	The mkgmir command starts Global Mirror for a specified session.
mkhostconnect	I/O port and host connect configuration	The mkhostconnect command configures open systems hosts port attachments to fibre-channel ports that are configured for FC-AL or SCSI-FCP topology. Open systems hosts port attachments to fibre-channel ports are configured for identified access mode and SCSI protocol.
mklcu	storage configuration	The mklcu command creates a logical control unit (LCU) in a storage image.
mkpprc	Copy Services	The mkpprc command establishes a remote mirror and copy (formerly PPRC) relationship for a volume pair.
mkpprcpath	Copy Services	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, 1750:1750, 1750:2105.
mkrank	storage configuration	The mkrank command creates one fixed block or count key data (CKD) rank from one array.
mkremoteflash	Copy Services	The mkremoteflash command initiates a remote point-in-time copy from source volumes to target volumes through a Remote Mirror and Copy relationship.
mksession	Copy Services	The mksession command opens a Global Mirror session.
mkuser	user account and security	The mkuser command creates a DS CLI or a DS Storage Manager user account. A CLI user with administrative authority uses this command to create a user account with a password and user group authority.
mkvolgrp	storage configuration	The mkvolgrp command creates a volume group in a storage image.

Command	Type	Description
offloadauditlog	audit commands	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.
pausegmir	Copy Services	The pausegmir command pauses Global Mirror for the specified session.
pausepprc	Copy Services	The pausepprc command pauses an existing remote mirror and copy volume pair relationship. Or, this command can be used to pause a single volume ID. To use 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.
resumegmir	Copy Services	The resumegmir command resumes Global Mirror processing for a specified session.
resumepprc	Copy Services	The resumepprc command resumes a remote mirror and copy (formerly PPRC) relationship for a volume pair.
resyncflash	Copy Services	The resyncflash command is a point in time copy of an existing FlashCopy pair 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. 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. When the resyncflash command is issued on the pair, the new data that is written to the source is copied to the target. The parameters specified in this command replace the parameters in the existing relationship. In order to keep the initial -record and -persist parameters, the -record and -persist parameters must be specified in the resyncflash command.

Command	Type	Description
resyncremoteflash	Copy Services	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.
reverseflash	Copy Services	The reverseflash command reverses the FlashCopy relationship.
revertflash	Copy Services	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.
revertremoteflash	Copy Services	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.
rmarray	storage configuration	The rmarray command deletes arrays.
rmckdvol	storage configuration	The rmckdvol command deletes count key data (CKD) base or alias volumes from a storage image.
rmextpool	storage configuration	The rmextpool command deletes extent pools from a storage image.
rmfbvol	storage configuration	The rmfbvol command deletes fixed block volumes from a storage image.
rmflash	Copy Services	The rmflash command removes a relationship between FlashCopy volume pairs.
rmgmir	Copy Services	The rmgmir command ends Global Mirror processing for the specified session.
rmhostconnect	I/O port and host connect configuration	The rmhostconnect command removes a SCSI host port connection from a storage image.
rmlcu	storage configuration	The rmlcu command deletes existing logical control units.

Command	Type	Description
rmpprc	Copy Services	The rmpprc command removes a Remote Mirror and Copy volume pair relationship. Or, this command can be used to remove 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). To use 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. The -unconditional parameter must be specified when you designate a single volume; otherwise an error occurs and the command process fails.
rmpprcpath	Copy Services	The rmpprcpath deletes a Remote Mirror and Copy path.
rmrank	storage configuration	The rmrank command deletes ranks from a storage image.
rmremoteflash	Copy Services	The rmremoteflash command removes a relationship between remote FlashCopy volume pairs.
rmsession	Copy Services	The rmsession command closes an existing Global Mirror session.
rmuser	user account and security	The rmuser command removes a storage image user account. CLI users with administrative authority use this command to delete a user account file. Administrators use their passwords in the required field.
rmvolgrp	storage configuration	The rmvolgrp command deletes existing volume groups from a storage image.
setflashrevertible	Copy Services	The setflashrevertible command modifies a FlashCopy volume pair that is part of a Global Mirror relationship to <i>revertible</i> . The revertible feature allows data to be committed to the target to form a new consistency, or reverted back to the last consistency. This command must be run before the FlashCopy pair can be committed or reverted.
setioport	I/O port and host connect configuration	The setioport command configures one or more I/O ports for open systems or zSeries host system connections. This command cannot be used for ESCON ports.

Command	Type	Description
setnetworkport	network ports	The setnetworkport command (DS8000 only) configures one network port for TCP/IP connections.
setremoteflashrevertible	Copy Services	The setremoteflashrevertible command modifies a remote FlashCopy volume pair that is part of a Global Mirror relationship to <i>revertible</i> . This command must be run before the FlashCopy pair can be committed or reverted. Once a pair is revertible, the data can be committed to the target to form a new consistency, or reverted back to the last consistency.
setrmpw	User accounts and security	The setrmpw command changes the IBM TotalStorage Productivity Center Replication Manager password. Only a person that has administrator authority can invoke this command.
setvpn	storage complex configuration	The setvpn command starts an outbound virtual private network connection.
showarray	storage configuration	The showarray command displays detailed properties of a specific array.
showarraysite	storage configuration	The showarraysite command displays detailed properties of a specific storage image array site.
showckdvol	storage configuration	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.
showextpool	storage configuration	The showextpool command displays detailed properties or performance metrics of an extent pool.
showfbvol	storage configuration	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.
showgmir	Copy Services	The showgmir command displays detailed properties and performance metrics for a Global Mirror logical subsystem ID.
showgmircg	Copy Services	The showgmircg command displays consistency group status for the specified Global Mirror session.

Command	Type	Description
showgmiroos	Copy Services	The showgmiroos command displays the number of unsynchronized (out of sync) tracks for the specified Global Mirror session.
showhostconnect	I/O port and host connect configuration	The showhostconnect command displays detailed properties of a storage image host connection.
showioport	I/O port and host connect configuration	The showioport command displays properties of an I/O port. It optionally displays the performance metrics for a specific I/O port.
showlcu	storage configuration	The showlcu command displays the detailed properties of an individual logical control unit (LCU).
showlss	storage configuration	The showlss command displays detailed properties of a logical subsystem (LSS).
shownetworkport	network port	The shownetworkport command (DS8000 only) displays detailed properties of an individual network port ID.
showpass	user account and security	The showpass command lists the properties of passwords.
showrank	storage configuration	The showrank command displays detailed properties or performance metrics of a rank.
showsi	storage image configuration	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.
showsp	storage complex configuration	The showsp command (for DS8000) displays detailed properties of a storage complex. Detailed properties include the names, descriptions, and customer account names for the storage complex.
showsu	storage unit configuration	The showsu command displays detailed properties of an individual storage unit.

Command	Type	Description
showuser	user account and security	The showuser command displays storage image user account details. A CLI user with administrative authority uses this command to display the properties (group assignment, user account status and number of failed log ins) that is associated with a current user account name.
showvolgrp	storage configuration	The showvolgrp command displays detailed properties of a volume group.
unfreezeflash	Copy Services	The unfreezeflash command resets a FlashCopy consistency group that was previously established with the -freeze parameter when the mkflash or resyncflash commands were issued.
unfreezepprc	Copy Services	The unfreezepprc command resumes I/O activity on a storage unit where the freezepprc command has been issued. The unfreezepprc command resets the <i>queue full</i> condition for the primary volume. All queued writes to the source volume are written.
ver	application key and version	The ver command displays the versions of the command-line interface, Storage Manager, and licensed machine code.

Command flags

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

Flag	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 by pressing any key. Note: This flag can be used only with the ls type (for example, lsuser , lskey , lsserver) commands and the help (setoutput) command.
-r	<i>number</i>	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: This flag can be used only with the ls type (for example, lsuser , lskey , lsserver) commands and the help (setoutput) command.

Flag	Parameters	Description
-fmt	xml	Sets the output format to XML. Note: This option can be used only with list (for example, lsuser, lskey, lserver) commands
	stanza	Sets the output format to stanza. Note: This option can be used only with list (for example, lsuser, lskey, lserver) commands
	delim	Sets the output format to a table. You must set the column delimiter to a single character with the -delim flag. Note: This option can be used only with list (for example, lsuser, lskey, lserver) commands
	default	Sets the output to a space-separated plain text table. Note: This option can be used only with list (for example, lsuser, lskey, lserver) 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 valid character as a delimiter. Note: This option can be used only with list (for example, lsuser, lskey, lserver) commands
-hdr	on off	Turns the header on or off. The default is on.
-bnr	on off	Turns the banner on or off. The default is on.
-v	[on off]	Turns verbose mode on or off. The default is off. If you specify the -v flag and do not specify on or off, then verbose mode defaults to on.
-fullid	[on off]	Provides fully qualified IDs, which include the storage image ID, for every ID that is displayed in the command output. The default value is off. Note: This command flag can only be used with list (for example, lsioport, lskey) and show (for example, showsu, showlss) commands.

Using hexadecimal values

Use this conversion table to determine the hexadecimal value for a field or DS CLI command parameter that requires a hexadecimal value, or to determine the decimal value for a field or DS CLI command parameter that displays a hexadecimal value, such as the LSS and LUN fields or values.

Purpose

Table 4 provides a list of common decimal to hexadecimal conversion values. For values not listed in the table, you can use the Windows calculator to convert values. From the Windows calculator, click **View**, then **Scientific**. If you are converting from decimal to hexadecimal, ensure that **DEC** is selected in the row of radio buttons, enter a decimal number, and then select **HEX**. If you are converting from hexadecimal to decimal, ensure that **HEX** is selected in the row of radio buttons, enter a hexadecimal number, and then select **DEC**.

Table 4. Decimal to hexadecimal conversion

Decimal	Hexadecimal
1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	8
9	9
10	A
11	B
12	C
13	D
14	E
15	F
16	10
17	11
18	12
19	13
20	14
21	15
22	16
23	17
24	18
25	19
26	1A
27	1B

Table 4. Decimal to hexadecimal conversion (continued)

Decimal	Hexadecimal
28	1C
29	1D
30	1E
31	1F
32	20
64	40
100	64
128	80
200	C8
256	100
500	1F4
1000	3E8
8192	2000
10000	2710
64384	FB80
65280	FF00

Command equivalents

Use this list to correlate commands that are supported on the Enterprise Storage Server machine type 2105 to equivalent commands on the DS8000 machine type 2107.

Enterprise Storage Server machine type 2105 command	DS8000 machine type 2107 command	Description
list server	lsserver	Like the 2105, a 2107 storage image contains one pair of servers. A 2107 storage image can contain two storage images.
list volumespace	lsextpool, showextpool, lsrank, showrank, lsarray, showarray, showarraysite	See Note 1.
create volumespace	mkextpool, mkarray, mkrank	
delete volumespace	rmrank, rarray, rmextpool	
list diskgroup	lsarraysite, showarraysite	Like the 2105 disk group, a 2107 array site consists of eight storage devices that are made into a RAID array. The 2107 does not support the JBOD array configuration.

Enterprise Storage Server machine type 2105 command	DS8000 machine type 2107 command	Description
list port	lsioport, showioport	<p>Like 2105, the 2107 supports fibre-channel and ESCON ports. The 2107 does not support parallel SCSI ports.</p> <p>The maximum quantity of host device adapter cards and I/O ports is dependent on 2107 model number and on the quantity of installed I/O enclosure features.</p> <p>The 2107 CLI lsioport and showioport commands include the -metrics parameter that returns the performance counter values for the respective I/O port IDs. The -metrics parameter provides the means to monitor I/O port performance statistics.</p> <p>For 2107, a I/O adapter card is assigned to a storage image.</p>
set port	setioport	See Note 2.
list volume	lsfbvol, lsckdvol	See Note 3.
create volume	mkfbvol, mkckdvol	
set volume	chfbvol, chckdvol	
list pav	lsckdvol, showckdvol	
create pav	mkckdvol	
delete pav	rmckdvol	
list volumeaccess	lsvolgrp, showvolgrp	See Note 4.
create volumeaccess	mkvolgrp, chvolgrp	
delete volumeaccess	rmvolgrp	
list hostconnection	lshostconnect, showhostconnect	<p>The 2105 and 2107 CLI commands are essentially the same, except that the 2107 commands include the volume group ID parameter.</p> <p>For 2107, the hostconnect commands concern SCSI-FCP host port connections to ESS I/O ports that are configured for SCSI-FCP and identified access mode.</p>
create hostconnection	mkhostconnect	
delete hostconnection	rmhostconnect	
set hostconnection	chhostconnect, managehostconnect	
list log	Not applicable.	--
list featurecode	lsuser, mkuser, rmuser, chuser, lstgencl	The 2107 CLI commands can display feature codes when the appropriate parameters are used with the commands.
list webuseraccount	Not applicable.	--
create webuseraccount	Not applicable.	
set webuseraccount	Not applicable.	
delete webuseraccount	Not applicable.	

Enterprise Storage Server machine type 2105 command	DS8000 machine type 2107 command	Description
list perfstats	lsioport, showiport, showrank, showextpool, showfbvol, showckdvol	The 2105 CLI list perfstats commands concern the Specialist facility that streams performance counter device adapter to the ESS Expert at predefined intervals. This facility does not exist for 2107. Use the 2107 CLI commands with the -metrics parameter to obtain current performance counter values.
create perfstats	Not applicable.	--
delete perfstats	Not applicable.	--
show remotesupport	showsp	The 2105 Specialist remote communication functions are implemented in the 2107 as attributes of the storage complex object. The 2107 CLI chsp and showsp commands provide methods to view, set, and modify the remote communication attributes. The 2107 does not support the pager remote communication function.
set remotesupport	chsp	
show email	showsp	
create email	chsp	
delete email	chsp	
set email	chsp	
show pager	Not applicable.	
create pager	Not applicable.	
delete pager	Not applicable.	
set pager	Not applicable.	
show snmp	showsp	
create snmp	chsp	
set snmp	chsp	
delete snmp	chsp	
list problem	Not applicable.	--
show problem	Not applicable.	--
delete problem	Not applicable.	--
list task	Not applicable.	--
show task	Not applicable.	--
list pprcpaths	lsflash, lsremoteflash, lspprcpath, lspprc, lsavailpprcport, showgmir	Unlike the 2105, the 2107 CLI Copy Services functions are not task-oriented. The 2107 CLI provides a complete set of FlashCopy and remote mirror and copy (formerly PPRC) make, change, remove, list, and show commands.
rsExecuteTask	Copy Services commands	The 2107 CLI provides a complete set of FlashCopy, remote mirror and copy (formerly PPRC) commands that can be used in the coding of scripts that emulate 2105 Copy Services tasks.
rsList2105s	lshostvol	The lshostvol command displays the mapping of host device or volume names to 2107 and 2105 and volume IDs.
rsPrimeServer	Not applicable.	--

Enterprise Storage Server machine type 2105 command	DS8000 machine type 2107 command	Description
rsQuery, rsQueryComplete, rsFlashCopyQuery	lsflash, lsprrc	<p>These 2107 Copy Services CLI commands are equivalent to the respective 2105 CLI commands.</p> <p>The 2107 mkflash and mkpprc commands provide a -wait flag that delays command response until copy complete status is achieved.</p>
rsTestConnection	ver	--
<p>Note 1</p> <p>Volume space configuration is a primary difference between 2105 and 2107. For 2105, one command configures an array site into a RAID array and rank. For 2107, one command configures an array site into an array, and a second command configures an array into a rank. For 2105, a rank is configured as fixed block or CKD, and a CKD rank can contain “interleave” CKD volumes. For 2107, a rank is assigned to a user-defined extent pool object, which the user defines as either the fixed block or CKD storage type. The “interleave” volume construct does not exist for 2107. For 2105, a volume is configured from a specific rank, and cannot span rank boundaries. For 2107, a volume is configured from an extent pool. An extent pool can contain multiple ranks. A 2107 volume consists of one or more extents that can be allocated from one or more ranks. A fixed block extent is 1 GB (128 logical blocks). Each block contains 512 bytes of usable data space. A CKD extent is 0.94 GB or 1113 CKD cylinders.</p> <p>For 2105, a rank is either assigned to server 0 or server 1, dependent on array site location. A 2105 rank is assigned to one of four possible LSS IDs, dependent on device adapter pair location and storage type configuration.</p> <p>For 2107, an extent pool is assigned to server 0 or server 1. A rank that is configured from any array site can be assigned to a server 0 or 1 extent pool. Array site position and device adapter pairs are not factors for the rank-to-extent-pool assignment. A volume that is created from a server 0 extent pool is assigned to an even-numbered LSS ID. A volume created from a server 1 extent pool is assigned to odd-numbered LSS ID. A user must define at least two extent pools (0 and 1) but can define as many extent pools as there are ranks. For 2105, a user can delete a rank but cannot delete a volume. For 2107, a user can delete a single volume, rank, or extent pool. The 2107 CLI showrank and showextpool commands include a -metrics parameter that returns the performance counter values for a specified rank or extent pool ID. The -metrics parameter provides the means to monitor rank and extent pool performance statistics.</p>		

Enterprise Storage Server machine type 2105 command	DS8000 machine type 2107 command	Description
<p>Note 2</p> <p>A 2107 ESCON I/O port is used for zSeries host attachment but cannot be configured as a remote mirror and copy path. Each ESCON I/O port must be assigned to only one address group. An address group is a set of 16 LSS IDs that are configured as CKD LCUs; for example, 0x00 to 0x0F. ESCON I/O port access to 2107 CKD volumes is constrained to the address group LCU volumes, up to 4096 volumes.</p> <p>A 2107 fibre-channel port is configured for either SCSI-FCP or FICON protocol. Like 2105, a FICON port is restricted to the point-to-point/switched fabric topology setting. A FICON I/O port is used for zSeries host attachment, but cannot be configured as a remote mirror and copy path. A FICON port must be configured for “anonymous” access mode, meaning that any zSeries host system port (WWNN or WWPN) has unrestricted access to all CKD volumes, up to 64 KB volumes.</p> <p>Like 2105, a 2107 fibre-channel SCSI-FCP I/O port can be configured for either the point-to-point/switched fabric or FC-AL connection topologies. A port that uses the point-to-point/switched fabric topology can be simultaneously used for OS host system I/O and for remote mirror and copy path configurations. Like 2105, a 2107 fibre-channel SCSI-FCP I/O port allows only “identified” host system ports to access volumes. A host system port WWPN must be identified (registered) to each I/O port through which volume access is intended. For 2107, this configuration constraint is defined as I/O port “identified” access mode. Host system port WWPN identification is accomplished by the CLI mkhostconnect command.</p>		
<p>Note 3</p> <p>A 2107 storage image can contain up to 32 000 volumes, whereas a 2105 unit can contain up to 8 000 volumes. Otherwise, the 2105 and 2107 volume definitions and characteristics are essentially identical.</p> <p>For 2107 CKD PAV volumes, the CLI list and show commands identify both the original base and current base volume assignments. The original and current base concept exists for 2105, but specific relationships are not identified in the output.</p> <p>The 2107 CLI provides a specific set of volume commands for each storage type (fixed block or CKD) as a means to clarify input parameter and output device adapter definitions.</p> <p>The 2107 CLI showfbvol and showckdvol commands include a -metrics parameter that returns the performance counter values for a specified volume ID. The -metrics parameter provides the means to monitor volume performance statistics.</p>		

Enterprise Storage Server machine type 2105 command	DS8000 machine type 2107 command	Description
Note 4		
<p>The 2105 volume access commands concern volume ID assignment to a SCSI-FCP host port initiator or WWPN. For 2107, volume IDs are assigned to a user-defined volume group ID (mkvolgrp and chvolgrp). A volume group ID is then assigned to one or more host system ports (mkhostconnect and chhostconnect) as a means to complete the volume access configuration.</p> <p>The volume group construct also exists in the 2105 internal code, but the construct is not externalized by the 2105 Specialist or CLI commands.</p> <p>For 2107, a user must create a FICON/ESCON-all type volume group. This volume group ID is assigned to each ESCON I/O port and to each FICON I/O port. The volume group ID enables FICON access all storage image CKD volumes, up to 64 KB volumes. The volume group ID enables an ESCON I/O port to access to the storage image CKD Address Group volumes, up to 4 KB volumes.</p> <p>For 2107 fixed block volumes, a volume group must be configured as either "SCSI-mask" or "SCSI-map-256", depending whether the volume group is accessed by a SCSI-FCP host port that uses the report LUNs or poll LUNs access method protocol.</p>		

Output field descriptions

This list describes the output field names, abbreviations, and field descriptions.

Output field name	Abbreviation	Description
Access	access	Access state. The term "access" is used throughout the CLI in different contexts.
Access state	accstate	The accessibility state of an object, online or offline.
Account	acct	Customer account name for a storage complex.
Active copy	actcpy	The FlashCopy background copy process is active.
Addr group	addrgrp	A set of 16 contiguous logical subsystems (LSSs) or logical control units (LCUs), starting at ID X0. Address group identifier that is assigned to this ESCON I/O port.
Address group	addrgrp	A set of 16 contiguous logical subsystems (LSSs) or logical control units (LCUs), starting at ID X0. Address group identifier that is assigned to this ESCON I/O port.
Array	array	An array of device drive modules (DDMs).
Array ID	arrayID	Array identifier that is assigned to a rank.

Output field name	Abbreviation	Description
Array site	arsite	Storage unit identifier followed by an array site identifier. Array site ID does not indicate physical location.
Attaching topology	atchtopo	Ports to which the host can attach.
Available storage	availstor	Storage that is available for a segment pool, in GB.
Background copy	bkgndcopy	FlashCopy process that copies data from a source volume to a target volume.
Base Vol #	basevolnum	Lowest number of logical volumes in the address group.
Bypass cache	bypasscach	Count of bypass cache I/O requests.
Bytes read	byteread	Count of bytes that are transferred by SCSI read I/O operations, in 128 KB.
Bytes written	bytewrit	Count of bytes that are transferred by SCSI write I/O operations, in 128 KB.
Cache fast-write hits	cachfwhits	Count of cache fast-write write I/O operations where data did not move to or from a storage device.
Cache fast-write read hits	cachfwrhits	Count of cache fast-write read I/O operations where data was not moved to or from a storage device.
Cache fast-write read requests	cachfwrreqs	Count of cache fast-write read I/O operations issued by a host to a volume.
Cache fast-write requests	cachfwrreqs	Count of cache fast-write write I/O operations issued by a host to a volume.
Cache space delay	cachspdelay	Count of delayed I/O operations from a host to a volume because of insufficient cache space.
Capacity	cap	Quantity of volume logical blocks or cylinders that are available for access by a host system.
CG attempts	CGattem	Number of attempts to form a consistency group.
CG drain	CGdrain	The maximum time that writes are inhibited to the remote site before stopping the current consistency group.
CG interval	CGinterval	The interval time between attempts to form a consistency group.
CG success	CGsuccess	The percentage of successful attempts to form consistency groups.
CG time	CGtime	The time when the last successful consistency group was formed.

Output field name	Abbreviation	Description
CKD irregular track accesses	CKDirtrak	Count of I/O operations from a host to a CKD volume that has accessed at least one logical track not described by a regular track format descriptor.
CKD irregular track hits	CKDirtrkhits	Count of irregular track I/O operations where data did not move to or from a storage device.
CKD write promote hits	CKDwrtprohits	Count of write I/O operations to a volume, where track format descriptor data in cache has promoted at least one track to cache with requiring access to a storage device.
Config Vols	confgvol	Number of logical volumes configured on an address group.
Configuration	config	Storage unit internal I/O interface configuration: model and feature code dependent.
Configured volumes	confgvols	Number of configured volumes.
Contaminating writes	contamwrts	Count of side file additions to a volume from an update to a concurrent copy protected track or an update to an XRC monitored track.
Control ops received	conopsrzd	Count of remote mirror and copy SCSI control I/O operations that are received from a SCSI source.
Control ops sent	conopssent	Count of remote mirror and copy SCSI control I/O operations that are sent to a SCSI target.
Control unit base type	conbasetype	Default or user-assigned logical control unit type.
Copy state	cpystate	Global Mirror copy state.
CopyIndicator	CopyIndicator	Indicates Yes if the CopyIndicator is set for this FlashCopy relationship.
Created	created	Date of creation.
Crit mode	critmode	Status of critical heavy mode, either enabled or disabled.
Critical heavy mode	crithvmode	Status of critical heavy mode for remote mirror and copy (formerly PPRC) copy operations, either enabled or disabled.
Current base vol	curbasevol	Base and alias volume number.
Current FICON logins	curflogs	Current number of FICON N-port worldwide node name (WWNN) identifiers that are logged in to this I/O port.
Current SCSI logins	currentlogs	Current number of N-ports that are logged in to this I/O port.

Output field name	Abbreviation	Description
Current time	currtime	Current date, time, local time zone, and Daylight Savings Time.
DA pair	DA pair	Identifier of the device adapter pair that the DDM is associated with. The DA pair indicates I/O enclosure location.
DASD cache transfers	DASDtrans	Count of logical tracks for a volume that were promoted to cache in full or partial track mode (excluding sequential pre-state).
Data	data	Status of the array data access: normal, degraded, read only, failed, repairing, or inaccessible.
Data state	datastate	Status of data access: normal, degraded, read only, failed, repairing, or inaccessible.
Date	date	Date and time, including time zones and Daylight Savings Time.
Date (85)	date	Current clock setting of date.
DDM capacity (GB)	DDMcap	Minimum disk capacity of DDMs, in GBs.
DDM RPM (revolutions per minute)	DDMRPM	Minimum rate of disk revolutions per minute of the DDMs in an array.
Description	desc	User-defined description.
Disk capacity (GB)	dkcap	Capacity of the DDM, in GBs.
Disk interface	dkinf	Interface type of DDM.
Disk rate (Gb/Sec)	dkrate	Interface rate of DDM, in GBs per second.
Disk RPM	diskrpm	DDM revolutions per minute.
Disk usage	dkuse	DDM usage in an array site.
Dynamic relocation source	dyrelocsource	Counting number of extents that were source of a dynamic relocation.
Dynamic relocation target	dyreloctarget	Counting number of extents that were target of a dynamic relocation.
EC level	EClvl	Engineering change level of the listed enclosure.
E-mail addresses	emailaddr	One or more e-mail addresses that receive service notification.
Enclosure #	enclnum	Identifier for an I/O enclosure within a storage unit frame.
Ending track	endtrk	Ending track address or number.
ESS IO ports	ESSIOport	The set of Enterprise Storage Server I/O ports to which a SCSI host port can log in.
ESSNet	ESSnet	Status of storage complex ESSNet user interface, either enabled or disabled.

Output field name	Abbreviation	Description
Extended long busy active	xtndlbzactive	Default or user-assigned extended long busy setting, enabled or disabled.
Extended long busy timeout (secs)	xtndlbztimeout	Default or user-assigned extended long busy timeout value.
Extent limit	extlim	Maximum number of possible extents.
Extent number	extentnum	LSS persistent cache extent number.
Extent pool	extpool	Extent pool.
Extent pool ID	extpoolID	Identifier for the extent pool of the assigned rank.
Extent pool name	extpoolnam	Name of the extent pool of the assigned rank.
Extent size	extsize	Number of logical tracks in an extent.
Extent threshold	extthresh	Extent allocation threshold setting that triggers notification.
Extents	exts	Number of extents in the rank.
Extents used	extused	Number of extents used by this volume ID.
Fan speed	fanspeed	Current speed of a fan.
Fatal reason	fatalrsn	Reason code for a fatal error.
Feature code	FC	Identifier code that is used to order the PC enclosure.
FICON	FICON	FICON I/O operations that are enabled for this port.
FICON enabled	fenabled	FICON status, enabled or disabled.
Firmware level	firmwarelevel	Identifier for the firmware level that is installed in the hardware management console (HMC) enclosure.
Frame	frame	Identifier of storage unit frame that contains this I/O enclosure. Frame identifier format is six hexadecimal characters.
Frame ID	frameID	Identifier of storage unit frame that contains this I/O enclosure. Frame identifier format is six hexadecimal characters.
Frame #	frm#	Frame number of a listed enclosure.
Group	group	--
Host ID	hostID	SCSI host identifier for an Open Systems host that is associated with this host port.
ID	ID	Storage image ID in the following format: manufacturer.type.serial number.

Output field name	Abbreviation	Description
I'mACopy	imacpy	Indicates Yes if the I'mAcopyBit is set for this FlashCopy relationship.
Inhibit cache loading	inhcachload	Count of inhibit cache loading I/O requests.
Interface address	interadd	FlashCopy arbitrated-loop base address of the storage image enclosure.
Interface IDs	interfID	Identifies four interface IDs that are associated with I/O ports on the HBA.
Interface rate	interrate	Minimum disk interface rate of the disk in an array, in GBs per second.
Interface type	interface type	Host attachment interface type (FC-AL, SCSI-FCP, or FICON).
Interfaces	interfs	Identifier of three interface ports for the HMC enclosure consisting of four hexadecimal characters.
Location	loc	Location of enclosure. Location format is <i>Utttt.mmm.ppsssss</i> .
Logical block	logblk	Logical block (512 bytes or 520 bytes).
Logical block size	logblksz	Logical block size.
Logical vols	logvols	Identifier of logical volumes.
Login limit	loglim	Maximum number of N-ports that can log in to this I/O port.
LSSs	LSSs	Number of logical subsystems in an address group.
LUN access	LUNacc	Quantity of LUNs that are accessible by this host attachment (256 LUNs or 64K LUNs).
LUN capacity	LUNcap	Quantity of LUNs that are accessible by this host attachment (256 LUNs or 64K LUNs).
Name	name	User-defined name.
NVS space allocations	NVSspallo	Count of I/O operations that cause nonvolatile storage (NVS) space allocation.
NVS space delay	NVSspadel	Count of I/O operations from a host to a volume.
Master count	mastcount	Quantity of master LSS IDs on this storage image ID.
Master ID	mastID	Master storage image ID.
Master session ID	mastsessID	Global Mirror session ID.
Master SSID	mastSSID	LSS subsystem identifier.
Memory	mem	Amount of processor memory in this PC enclosure, in MB.
Migration disk SN	migradiskSN	Migration disk serial number.

Output field name	Abbreviation	Description
MRPD interval	MRPDintvl	Number of days between dial home to report machine-reported product data (MRPD). This interval can be 0.
MTS	MTS	Manufacturer, machine type, and sequence number.
Narrow arrays	nararrays	Number of narrow arrays in a rank.
Normal read hits	normrdhits	Number of normal read operations where data did not move to or from a storage device.
Normal read requests	normrdrqts	Number of normal read operations issued by a host to a volume.
Normal write hits	normwritehits	Number of normal write operations where data did not move to or from a storage device.
Normal write operations	normwriteops	Number of command sequences with at least one write command.
Normal write requests	normwritereq	Number of normal write operations issued by a host to a volume.
NVS memory	NVSmem	Amount of nonvolatile storage (NVS) memory that is assigned to this server, in MB.
Number of logins	numlogins	Current number of valid N-ports that are logged in to this I/O port.
Number of logical vols	numlvols	Number of logical volumes configured from an extent pool.
Number of extents	numexts	Number of extents.
Number of ranks	numranks	Number of ranks configured in an extent pool.
Original base vol	orgbvols	Original base volume CKD volume identifier.
OS memory	OSmem	Amount of operating system memory that is assigned to the server, in MB.
Out-of-sync tracks	outsynctrks	The number of tracks that are not synchronized for this FlashCopy relationship.
Persistent	persistent	Status of persistent FlashCopy, either enabled or disabled.
Persistent cache (MB)	pcache	Amount of persistent cache memory that is assigned to a server, in MB.
Physical bytes read	phbyteread	Number of physical bytes read.
Physical bytes written	phbytewrit	Number of physical bytes written.
Physical read	phread	Number of physical storage read operations.
Physical write	phwrite	Number of physical storage write operations.
Port	port	I/O port.
Port group	portgrp	Group identifier for host port.

Output field name	Abbreviation	Description
Port profile	portpro	Port behavior identification for this SCSI host ID.
Position	pos	Position of DDM in an array configuration of DDMs.
Power mode	pw mode	Current storage unit power control mode.
Power state	pw state	Current storage unit power status.
PPRC control operations received	PPRCcntroprec	Remote mirror and copy (formerly PPRC) I/O control operations.
PPRC control operations sent	PPRCcntropsent	Remote mirror and copy (formerly PPRC) I/O control operations.
PPRC source	PPRCsce	Remote mirror and copy (formerly PPRC) source I/O operations that are enabled for this port.
PPRC target	PPRCtgt	Remote mirror and copy (formerly PPRC) target I/O operations that are enabled for this port.
PPRC tracks	PPRCtrks	Count of logical tracks for a remote mirror and copy (formerly PPRC) primary volume that were transferred to a remote mirror and copy secondary volume.
Processor complex	procplex	Identifier of processor complex with which the central electronic complex enclosure is associated.
Processor qty	procqty	Number of processors in the PC enclosure.
Processors assigned	procassd	Number of processors that are assigned to the server.
Profile	profile	Host port profile.
Quick write promotes	qwritepromts	Count of logical tracks for a volume that have been destaged from cache to storage devices.
RAID type	RAIDtype	Type and configuration of RAID array.
Rank	rank	Identifier that the array is assigned to. Rank number consists of storage unit ID and a rank number.
Rank group	rnkgrp	Identifier of rank group where segment pool is configured.
Rank position	rankpos	Array position within the assigned rank.
Read operations	readops	Count of I/O command sequences in one read or search command (but no write commands).
Reads	reads	Count of read I/O operations.
Real allocated extents	realallocext	Count of real allocated extents.
Real extent conversion	realextconv	Count of real extent conversion.

Output field name	Abbreviation	Description
Real extent pool capacity	relextcap	Number of gigabytes of real extent pool capacity.
Real extents	realext	Count of real extents.
Reason	reason	The reason a condition exists.
Record cache misses	reccachemis	Number of normal record mode read operations where data moved to or from a storage device.
Record mode reads	recmoreads	Number of normal record mode read operations issued by a host to a volume.
Recording	rec	FlashCopy - record changed tracks.
Requested capacity	reqcap	Number of volume cylinders that are available for host system access.
Requested power mode	reqpm	Power control mode when the local/remote switch is set to remote.
Requested state	reqstate	Desired state of storage unit: online or offline.
Reserved extents	resvdexts	Extents that are reserved in an extent pool.
Reserved storage (GB)	resvdstor	Reserved storage in an extent pool, in GB.
Revertible	revertible	Revertible to previous Global Mirror state.
SCSI host port users	SCSIhostportusr	SCSI host port users of this volume group.
SCSI TGT	SCSItgt	SCSI target I/O operations that are enabled for this port.
Scheduled on	schl-on	User-defined time that the storage unit powers on.
Scheduled off	schl-off	User-defined time that the storage unit powers off.
Sequential DASD cache transfers	seqDASDtrans	Count of logical tracks for a volume that were promoted to cache because data was moved by sequential pre-stage and required movement from a storage device.
Sequential read hits	seqreadhits	Number of sequential read operations where data did not move to or from a storage device.
Sequential read requests	seqreadreqs	Number of sequential read operations issued by a host to a volume.
Sequential write hits	seqwritehits	Number of sequential write operations that did not require movement of data to or from a storage device before the completion of the operation.
Sequential write operations	seqwriteops	Number of command sequences that contain at least one sequential write command.

Output field name	Abbreviation	Description
Sequential write requests	seqwritereq	Number of sequential write operations issued by a host to a volume.
Serial number	SN	Internal identifier for the data space of an array.
Server	server	Server or DA group to which the DA is assigned.
SN	SN	Unique serial number.
SNMP addresses	SNMPaddr	One or two IP addresses where the storage complex sends SNMP error messages.
Source write enabled	sourcewrite	Host writes to the source volume are allowed.
Speed	speed	The current speed of this fan tray.
SRC cascade	SRCcascade	Source volume is enabled to be in a cascading remote mirror and copy (formerly PPRC) relationship.
SRC vol LSS	SRCvolLSS	Source volume LSS.
SS	SS	Subsystem.
Starting track	starttrk	The starting track address for the volume pinned data.
State	state	Storage unit functional status: online, offline, resuming, quiescing, quiesce exception, forced quiescing, or fenced.
State (FRU)	statefru	Current state of the disk drive module.
Storage devices	stordev	Number of storage devices in an enclosure.
Storage slots	storslot	Number of slots for storage devices in an enclosure.
Storage type	stortype	Extent pool type of the assigned rank.
Storage unit	su	One storage device.
Strip size	strpsize	Number of logical tracks in a strip.
Stripe size	strpsize	Number of logical tracks in a stripe.
Subordinate count	subcount	Count of subordinate associations.
Subordinate ID	subID	Subordinate storage unit ID.
Subordinate SSID	subSSID	Subordinate subsystem identifier.
Subsystem	subsys	User-assigned or default subsystem identifier.
Suspended	suspended	The relationship is suspended.
Synced	synced	Date the FlashCopy was synchronized.
Target write enabled	tgtwrite	Host write I/O operations to a target volume are allowed.

Output field name	Abbreviation	Description
Tgt cascade	tgtcascade	Target volume is enabled to be configured in a cascading remote mirror and copy relationship.
Tgt read enabled	tgthead	Host read I/O operations to a target volume are allowed.
Time	time	Current clock setting of time.
Time lower interface activity	timelowifact	Accumulated time of lower interface I/O activity for the volume.
Time on-channel	timeonchan	Amount of I/O port time for SCSI I/O operations, in seconds.
Time physical storage read	timephread	Accumulated time for physical storage read operations.
Time physical storage write	timephwrite	Accumulated time for physical storage write operations.
Time read	timeread	Accumulated time for all read operations.
Time write	timewrite	Accumulated time for all write operations.
Time zone	timez	Current clock settings of time zone and Daylight Savings Time.
Timeout active copy recording	timoutactcpyrec	--
Total storage (GB)	totlstor	Amount of storage in an extent pool, in GB.
Topology	topo	Port topology.
Track size	trksize	Size of the track if the volume is CKD or fixed block.
Type	type	Type of storage unit enclosure.
Unknown SCSI IDs	unkSCSIlog	List of unknown SCSI N-port WWPN identifiers that attempted login into this I/O port.
Used extents	usedexts	Number of extents that are allocated to volumes in a rank.
Virtual extent conversion	virextconv	Count of virtual extent conversion.
Virtual extent pool capacity	virextcap	Number of gigabytes of virtual extent pool capacity.
Virtual extents	virext	Count of virtual extents.
Vol group	vol	Volume group ID. The unique identifier that is assigned to this volume.
Volume group	volgrp	Volume group.
Vols	vols	Number of logical volumes in an address group.
Volume type	voltype	Volume type.
Wide arrays	widearrays	Number of wide arrays in a rank.
Writes	writes	Count of write I/O operations.

Output field name	Abbreviation	Description
WWNN	WWNN	Worldwide node name.
WWPN	WWPN	Worldwide port name.
XDC interval	XDCintvl	Global Mirror copy interval.

Framework command

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

The framework commands consist of the following commands:

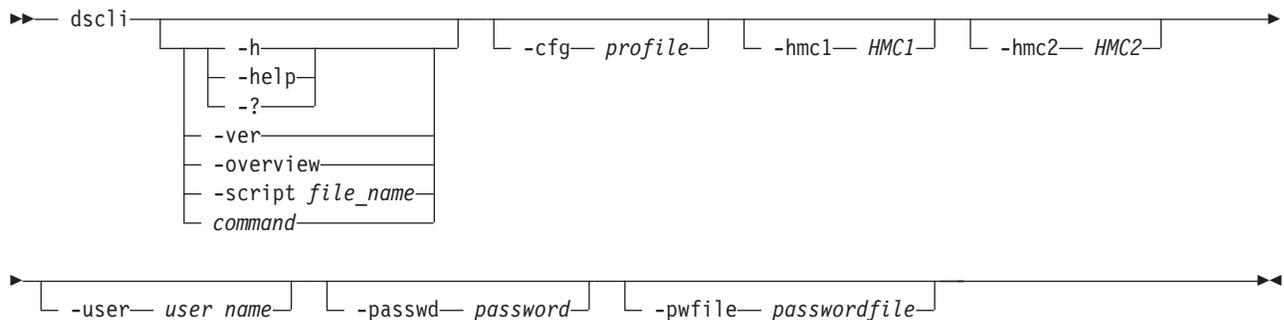
- **dscli**
- **exit**
- **help**
- **quit**
- **setoutput**

The **dscli** command starts the DS command-line interface (CLI). Use this command to perform storage management tasks from the command-line.

The **setoutput** command generates the format of the reports for the list and show commands during an interactive command session.

dscli

The **dscli** command starts DS CLI. Use this command to run DS CLI commands in the interactive, single-shot, or script mode.



Parameters

Note: You must not start more than 100 DS CLI sessions simultaneously. Starting more than 100 DS CLI sessions simultaneously can result in connection problems.

-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.

-script *file_name*

(Optional) Initiates the script mode so that multiple dscli program commands can be issued consecutively using a saved file.

file_name

Specifies the file with the stored program commands to be executed.

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.

command

Specifies the single command that you want to run.

-cfg *profile*

Specifies a profile file. This parameter is not required if you are using 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.

-hmc1 *HMC1*

(Optional) Specifies the primary management console IP address or the host name.

HMC1

The IP address for the primary management console.

This parameter is not required if you have established this information as a profile variable.

-hmc2 *HMC2*

(Optional) Specifies the secondary management console IP address or the host name.

HMC2

The IP address for the secondary management console.

This parameter is not required if you have established this information as a profile variable.

Note: The *HMC1* and *HMC2* values must apply to two different management consoles.

-user *user_name*

(Optional) Specifies your user name for issuing DS CLI commands on the command-line.

user_name

Your user name.

This parameter is not required if you have established this information as a profile variable.

-passwd *password*

(Optional and not recommended) Specifies the password that you use for issuing DS CLI commands on the command line.

password

Your password.

This parameter is **not** required or **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 application.

Also, if you specify this parameter and do not specify the **-user** parameter, nothing happens. In other words, you are still prompted for a user ID and password before you can log in to the DS CLI application.

-pwfile *passwordfile*

Specifies the password file that contains your password.

passwordfile

Specifies a password file as an alternative to the **-passwd** parameter.

Example

This command invokes the CLI in interactive mode:

```
>dscli
```

The resulting output

```
dscli>
```

exit

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.

Example

Exit interactive mode

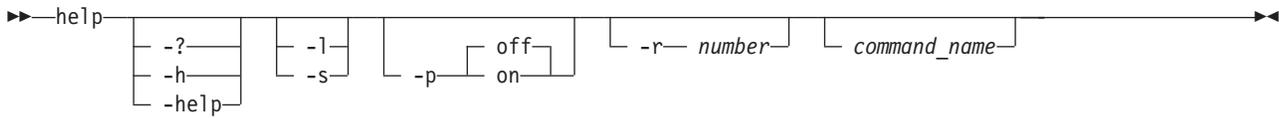
```
dscli>exit
```

Out of interactive mode

```
#
```

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.



Parameters

-? | **-h** | **-help**

Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

-l Displays a list of available commands with the syntax diagrams for each. If you specify a command name with this parameter, this command displays the syntax for only the specified command.

-s Displays a list of available commands with a brief description of each. If you specify a command name with this parameter, this command displays a brief description for only the specified command.

-p *off* | *on*

Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value.

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.

Example

Invoke help

```
#dscli>help -s exit
```

The resulting output

Ends a command-line interface session.

quit

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.

Example

Quit interactive mode

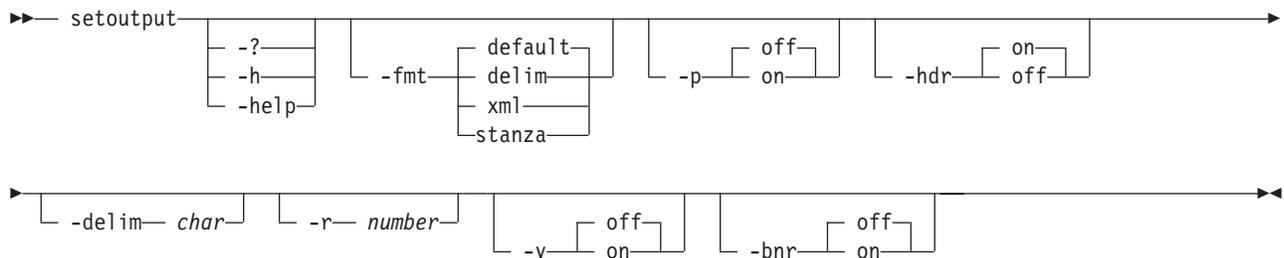
```
dscli>quit
```

Out of interactive mode

```
#
```

setoutput

The **setoutput** command sets or displays command output format options. You can use this command to set either default or user-defined output formats. The output format set by this command remains in effect for the duration of the interactive command session unless reset either with a command option or the re-issuance of the **setoutput** command. Running this command with no parameters displays the current output settings in the default output format.



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.

-fmt default | **delim** | **xml** | **stanza**

Specifies the format of the output. You can specify one of the following values:

default

Specifies that the output be displayed in a tabular format using spaces as the delimiter between the columns. This is the default value.

delim

Specifies that the output format be set to a table and sets the column delimiter to a single character specified by the **-delim char** option.

xml

Specifies that the output be displayed in XML format.

stanza

Specifies that the output be displayed in stanza (horizontal table) format.

- 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 is the default value.
 - on** Displays one page of text at a time. Pressing any key displays the next page.
- hdr** *on | off*
Specifies whether to display the table header.
 - on** Displays the table header. This is the default value.
 - off**
Does not display the table header.
- delim** *char*
Specifies the delimiter character (such as a comma) used in the report.
- 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.
- v** *off | on*
Specifies whether to enable verbose mode.
 - off**
Disables verbose mode. This is the default value.
 - on** Enables verbose mode.
- bnr** *off | on*
Specifies whether the banner (command header) message be enabled.
 - off**
Turns the header mode off so that the command header does not display.
 - on** Turns the header mode on so that the command header is displayed.

Format Examples

Invoke the **setoutput** command with no options

When you specify the **setoutput** command with no options, the DS CLI always displays the current output settings in the default format (space-separated plain text table), regardless of the values of the output settings. Issue the **setoutput** command as follows:

```
dscli>setoutput
```

The resulting output

```
Paging  Rows   Format  Header  Verbose  Banner
=====
Off     -    Default On  Off  On
```

Invoke the **setoutput** command using the **-delim** parameter

The following is an example of the commands that you would issue to get (long) output in comma-separated format for an unassigned rank only. Issue the **setoutput** command to specify the report format and then issue the **lsrank** command to designate the rank being queried.

```
dscli> setoutput -fmt delim -delim ,
dscli> lsrank -dev IBM.2107-75FA120 -state unassigned
```

The resulting output

Note: While this example shows the header turned on, you can choose to turn the header off, in which case you 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 is an example of the commands that you would issue to get (long) output in XML format for an unassigned rank only. Issue the **setoutput** command to specify the report format and then issue the **lsrank** command to designate the unassigned rank being queried.

```
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
</PROPERTY><PROPERTY NAME="raidtype"><DISPLAY TYPE="unit8">5</DISPLAY>
<VALUE TYPE="string">5</VALUE></PROPERTY><PROPERTY NAME="extpool_id">
<DISPLAY TYPE="string">-</DISPLAY><VALUE TYPE="string">-</VALUE>
</PROPERTY><PROPERTY NAME="stgtype"><DISPLAY TYPE="string">fb</DISPLAY>
<VALUE TYPE="string">fb</VALUE></PROPERTY>
</INSTANCE></IRETURNVALUE>
```

Invoke the setoutput command using the -fmt stanza parameter

When columns are horizontally long, output can be difficult to visually align. Using the stanza format option eliminates this problem. The following is an example of the commands that you would issue to get (long) output in stanza format for an unassigned rank only. Issue the **setoutput** command to specify the report format and then issue the **lsrank** command to designate the unassigned rank being queried.

```
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
```

User account and security commands

This section contains commands that are used to maintain command-line interface (CLI) user accounts and security.

The following commands are used to add, modify, delete, and show CLI user accounts and security information:

- **chpass**
- **chuser**
- **lsuser**
- **managepwfile**
- **mkuser**
- **rmuser**
- **showpass**
- **showuser**

The **chpass** command changes the password expiration time and the number of login attempts for a storage complex.

The **chuser** command modifies and locks or unlocks a DS CLI or a DS Storage Manager user account. Users that 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.

The **lsuser** command generates a report that lists the storage image user account names and access authority levels.

The **managepwfile** command creates a password file for an existing ESS or DS user account. This command processes the password requirements for 2105, 2107, and 1750 systems.

The **mkuser** command creates a DS CLI or a DS Storage Manager user account.

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

The **showpass** command generates a report that lists the number of days until the password expires and the number of failed logins that are associated with the password.

The **showuser** command generates a report that displays the details for an individual storage image user account.

chpass

The **chpass** command changes the password expiration time and the number of login attempts for a storage complex.

```
►► chpass [ -expire— number ] [ -fail— number ]
```

Parameters

-expire *number*

(Optional) Specifies the number of days a Storage Manager user account password is valid before it expires. The default number of days is 365. 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.

-fail *number*

(Optional) Specifies the number of login attempts allowed on any given Storage Manager user account. The number of login attempts can be zero to twenty-five. The default number of login attempts is 15. 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.

Example

Invoking the **chpass** command

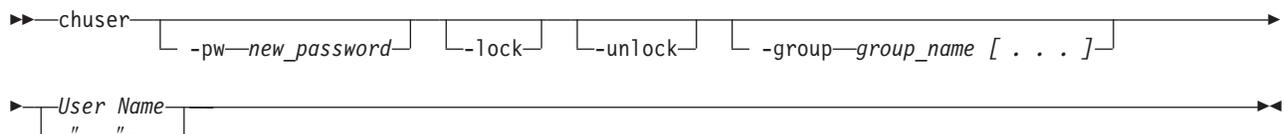
```
dscli>chpass -expire 20 -fail 0
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI  
Version 5.0.0.0 DS:IBM 2107-75FA120  
Password parameters successfully changed.
```

chuser

The **chuser** command is used to modify and lock or unlock a DS CLI or a DS Storage Manager user account. A CLI user with administrative authority uses this command to update a user account password, modify user group authority, or to lock or unlock a user account. Users that 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.



Parameters

Note: When a person with administrator authority designates the password, the password is set as expired upon its initial use. The user of the password is required to establish a new password using the **chuser** command before access to the rest of the DS CLI application is granted.

-pw *new_password*

(Optional) Specifies that a new password be assigned to the user.

Notes:

1. When a person with administrator authority is using 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 their valid password and replaces their prior password.

new_password

The new password.

The new password and its usage must meet the following criteria:

- Be six to 16 characters long.
- Must contain five or more letters, and it must begin and end with a letter.
- Must contain one or more numbers.
- Cannot contain the user's user ID.
- Is case-sensitive.

-lock

(Optional) Locks a user account.

Persons with administrator authority can use this parameter to lock a user account. The affect of this locking action is not enacted until the user authenticates their account. In other words, if a user is already active (authenticated) and using the DS CLI application, the lock does not take place until they log out of the application.

-unlock

(Optional) Unlocks a user account.

A person with administrator authority can use this parameter to unlock a user account when the user can no longer log into the DS CLI application. The reasons a user might not be able to log into the DS CLI application can include:

- The user forgot their password and in an attempt to log in they 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) The user's access authority group.

group_name [...]

The following list provides the list choices that can be assigned to a user. Multiple names must be separated by commas. For example, *op_copy_services,service*.

admin

The administrator user group has access to all management console server service methods and all storage image resources.

op_storage

The storage operator user group has access to physical configuration service methods and resources, including storage complex, storage image, array, rank, and extent pool objects. This user group inherits all the authority of the *op_copy_services* and *monitor* user groups, excluding security methods.

op_volume

The volume operator user group has 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 group inherits all authority of the monitor user group.

op_copy_services

The copy services operator user group has access to all Copy Services service methods and resources, excluding security methods. In addition, this user group inherits all authority of the monitor user group.

service

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

monitor

The monitor user group has access to list and show commands. It provides access to all read-only, nonsecurity management console server service methods and resources.

no access

The no access user group does not have access to any service methods or storage image resources. By default, this user group is assigned to any user account in the security repository that is not associated with any other user group.

User Name | -
(Required) The user account name.

Notes:

1. The administrator inserts the name of the user account that is affected by the changes (group name, lock, or unlocking).
2. Users who are changing their passwords insert their user account name.

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the chuser command

```
dscli>chuser -pw xy0abcde testuser
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
User Name testuser successfully modified.
```

lsuser

The **lsuser** command returns a list of storage image user account names and access authority levels.

▶▶—lsuser—▶▶

Parameters

There are no parameters for this command.

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 tables.

Invoking the lsuser command

```
dscli>lsuser
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI
```

Name	Group	State
Testuser	services,op_copy_services	active
Biguser	admin	active
Smalluser	op_storage	locked

Report field definitions

Name Specifies the user name that is assigned to the user account.

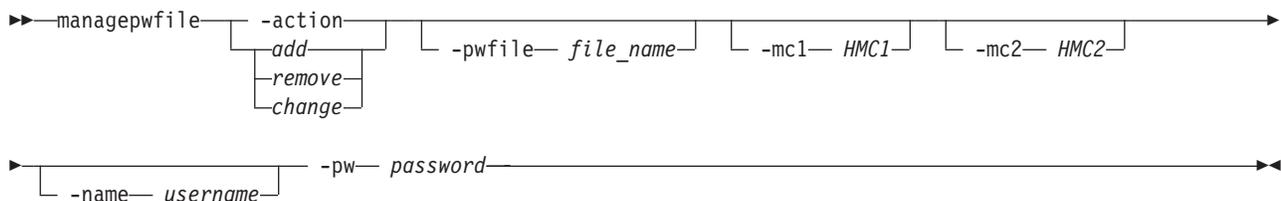
Group Specifies the access authority group of the user. One or more of the following group designations is displayed:

- admin
- op_storage
- op_volume
- op_copy_services
- service
- monitor
- no access

State Specifies the status of the user account for the designated user group, either active or locked.

managepwfile

The **managepwfile** command creates a password file for an existing ESS or DS user account. This command processes the password requirements for 2105, 2107, and 1750 systems.



Parameters

-action

(Required) Specifies that a process that is designated by the subparameter be enacted on the password file.

add Specifies that the password file be created by adding the first line of the file, which contains the primary HMC designation and the HMC user name.

remove Specifies that the password file be removed for the designated user.

change Specifies that the password file 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.

file_name
The user-specified ESS or DS password file name.

-mc1

(Optional) Specifies the DNS name or the IP address.

Note: You can use the default value that is specified for HMC1 in your profile file, or you can use the value that is specified for the current CLI session connection.

HMC1
Designates the primary HMC or 2105 DNS name or IP address. This information along with the **-name** parameter is used as a key in the password file.

-mc2

(Optional) Specifies the DNS name or the IP address of the secondary HMC.

Note: You can use the default value that is specified for HMC2 in your profile file, or you can use the value that is specified for the current CLI session connection.

HMC2
Designates the secondary HMC or 2105 DNS name or IP address.

-name

(Optional) Specifies the name that you use to access the DS CLI application.

username
Designates the user-specified HMC or 2105 user name. This information along with the **-mc1** parameter information is used as a key in the password file.

-pw *password*

(Required) Specifies a user-assigned password.

password
Specifies the password that enables user name access to this ESS or DS CLI client installation. The password is case-sensitive.

Notes:

1. A password file is created with a user's default protection mask. The user needs to 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.
2. 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 operating system

For a Windows operating system, the property value defaults to the environment variable `%USERPROFILE%`. As a result, your personal profile is `C:\Documents and Settings\\dscli\security.dat`.

UNIX or Linux operating system

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

i5/OS

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

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 if you have an environment setting like the following on your local system:

```
_JAVA_OPTIONS=-Duser.home=...
```

Example (2107)

Invoking the managepwfile command

```
dscli>managepwfile -action add -mcl myess.ibm.com -name testuser -pw AB9cdefg
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI
Version 5.0.0.0 DS: IBM.2107-75FA120
Record myess.ibm.com/testuser successfully added to password file
c:\Documents and Settings\testuser\dscli\security.dat
```

mkuser

The **mkuser** command creates a DS CLI or a DS Storage Manager user account. A CLI user with administrative authority uses this command to create a user account with a password and user group authority.

```
►►—mkuser— -pw— password— -group—group_name [ . . . ]—User Name—►►
                               [ " _ " ]
```

Parameters

pw *password*

(Required) Specifies the password that is assigned to the user that allows them to access the use of the DS CLI command line function. This password is temporary and set to expire 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 application.

password

The password assigned by the administrator to user.

The password and its usage must meet the following criteria:

- Be six to 16 characters long.
- Must contain five or more letters, and it must begin and end with a letter.
- Must contain one or more numbers.
- Cannot contain the user's user ID.
- Is case-sensitive.

-group *group_name* [...]

(Required) The user's access authority group.

group_name [...]

The following list provides the list choices that can be assigned to a user. Multiple names must be separated by commas. For example, `op_copy_services,service`.

admin

The administrator user group has access to all management console server service methods and all storage image resources.

op_storage

The storage operator user group has access to physical configuration service methods and resources, including storage complex, storage image, array, rank, and extent pool objects. This user group inherits all the authority of the `op_copy_services` and `monitor` user groups, excluding security methods.

op_volume

The volume operator user group has 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 group inherits all authority of the `monitor` user group.

op_copy_services

The copy services operator user group has access to all Copy Services service methods and resources, excluding security methods. In addition, this user group inherits all authority of the `monitor` user group.

service

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

monitor

The `monitor` user group has access to list and show commands. It

provides access to all read-only, nonsecurity management console server service methods and resources.

no access

The no access user group does not have access to any service methods or storage image resources. By default, this user group is assigned to any user account in the security repository that is not associated with any other user group.

User Name | -

(Required) The current user account user name.

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the mkuser command

```
dscli>mkuser -pw AB9cdefg -group service,op_copy_services testuser
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI
Version 5.0.0.0 DS: IBM.2107-75FA120
User Name testuser successfully created.
```

rmuser

The **rmuser** command removes a storage image user account. CLI users with administrative authority use this command to delete a user account file. Administrators use their passwords in the required field.



Parameters

-quiet

(Optional) Turns off the confirmation prompt.

User Name | -

(Required) The current user account user name.

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the rmuser command

```
dscli>rmuser testuser
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version 5.0.0.0 DS: IBM.2107-75FA120
Are you sure you want to delete User Name testuser? y/n
Y
User Name testuser successfully deleted.
```

setmpw

The **setmpw** command changes the IBM TotalStorage Productivity Center Replication Manager password. Only a person that has administrator authority can invoke this command.

```
▶▶ setmpw -dev storage_image_ID [-server 0|-1|-both] [-rmpw RM_password]▶▶
```

Parameters

-dev *storage_image_ID*

(Required) Specifies the storage image ID, which consists of the values for manufacturer, machine type, and serial number.

-server *0 | 1 | both*

(Optional) Specifies the server where the password is changed. The default is both.

-rmpw *RM_password*

(Required) Specifies the new Replication Manager password. The password must meet the following criteria:

- The password must contain a minimum of 5 alphabetic characters and 1 numeric character.
- The password must begin and end with an alphabetic character.
- The alphabetic characters must be entered in uppercase and the password is case sensitive.
- The password cannot contain the user name of the account that it is for.

Example

Note: If you do not specify the **-rmpw** parameter on your command line, you are asked twice to supply the Replication Manager password after you have entered the **setmpw** command.

Invoking the setmpw command

```
dscli>setmpw -dev IBM.2107-75ALA2P  
-server 0 -rmpw CLI1TIME
```

The resulting output

```
Sun OCT 8 02:23:49 PST 2006 IBM DS CLI  
Version 5.0.0.0 DS:IBM 2107-75ALA2P  
RM password was successfully updated.
```

showpass

The **showpass** command lists the properties of passwords.

```
▶▶ showpass▶▶
```

Parameters

There are no parameters for this command.

Example (2107)

Invoking the showpass command

```
dscli>showpass
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI  
Version 5.0.0.0 DS: IBM.2107-75FA120  
Password Expiration (days) 33 days  
Failed Logins Allowed 5
```

showuser

The **showuser** command displays storage image user account details. A CLI user with administrative authority uses this command to display the properties (group assignment, user account status and number of failed log ins) that is associated with a current user account name.

```
▶▶ showuser     User_Name     ▶▶  
          [ " _ " ]
```

Parameters

User_Name | -

(Required) Specifies the user account name.

If you specify the dash (-), this parameter information is automatically supplied.

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 reports that are associated with the **showuser** command. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750. .

Invoking the showuser command

```
dscli>showuser testuser
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI  
Version: 5.0.0.0 DS: IBM.2107-75FA120
```

Column Header	Description
Name	Name of the user that you have queried.

Column Header	Description
Group	The user's access authority group. One or more of the following group designations is displayed: <ul style="list-style-type: none"> • admin • op_storage • op_volume • op_copy_services • service • monitor • no_access
State	The status of the user account for the specified user group, either active or locked.
FailedLogins	Count of login failures since last successful login for this user. This number resets to 0 with each successful login.

Application key and version commands

This section contains commands that are used to maintain application keys and to view the current CLI version.

Use the following commands to add and list application keys and to show the current command-line interface version:

- **applykey**
- **lskey**
- **ver**

The **applykey** command activates the licenses for your storage unit.

The **lskey** command generates a report that displays the type of licensed machine code activation keys that are installed and are available for use by the specified storage unit.

The **ver** command displays the versions of the command-line interface, Storage Manager, or licensed machine code either individually or collectively depending on how you enter the command.

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. The file that contains the LMC keys must be downloaded from an IBM Web site.

```

▶▶ applykey [-key key [...]] [-file file_name] [storage_image_ID]

```

Parameters

-key *key [...]*

(Optional) Specifies the LMC key. 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 will be imported. The ID includes manufacturer, type, and serial number.

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the **applykey** command

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

lskey

The **lskey** command displays the type of licensed machine code (LMC) activation keys that are installed and available for use by the storage unit.

►— 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 specify the dash (-), this parameter information is automatically supplied.

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 **lskey** command. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750. The only difference is the value input for the *storage_image_ID* parameter.

Invoking the **lskey** command

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

The resulting output

```
Sun Apr 09 02:23:49 PST 2004 IBM DS CLI
```

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)	25.1	CKD
Point in time copy (PTC)	On	All
Parallel access volumes (PAV)	On	CKD
HyperPAV	On	CKD

Report field definitions

Activation key

Specifies the type of LMC activation key that is activated for the storage image. One of the following values is displayed:

- Operating Environment
- Point-in-time copy
- Remote Mirror and Copy
- Metro Mirror
- Global Mirror
- Metro/Global Mirror
- Parallel access volumes
- Remote Mirror for z/OS
- IBM HyperPAV

Authorization Level (TB)

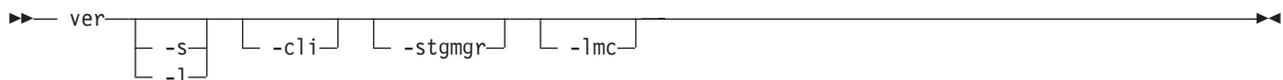
Specifies the capacity of the specified license feature. The quantity is displayed in terabytes (TB). One of the following values is displayed:

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

Scope Specifies 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.

ver

The **ver** command displays the versions of the command-line interface, Storage Manager, and licensed machine code.



Parameters

-s (Optional) The **-s** parameter displays the version of the command line interface program. You cannot use the **-s** and **-l** parameters together.

-l (Optional) The **-l** parameter displays the versions of the command line interface, Storage Manager, and licensed machine code. You cannot use the **-l** and **-s** parameters together.

-cli
(Optional) Displays the version of the command line interface program. Version numbers are in the format `version.release.modification.fixlevel`.

-stgmgr
(Optional) Displays the version of the Storage Manager.

-lmc
(Optional) Displays the version of the licensed machine code.

Example

Invoking the `ver` command

```
dscli>ver -l
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS
Storage Manager 5.0.1.0
LMC 5.0.1.0
```

Physical enclosure information commands

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

Use the following commands to view information about the physical enclosures in your storage complex:

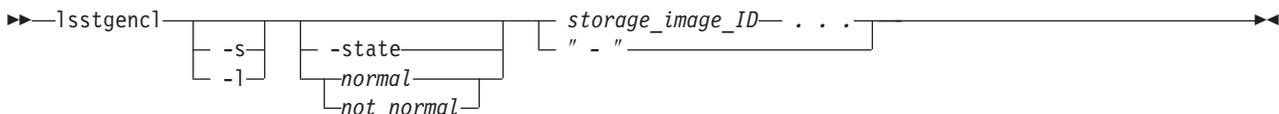
- **lsframe**
- **lsstgenc1**

The **lsframe** command generates a report that displays a list of the frame enclosures for a storage image.

The **lsstgenc1** command generates a report that displays a list of the storage enclosures and status information for each enclosure in the list.

lsstgenc1

The **lsstgenc1** command displays a list of storage enclosures and status information for each enclosure in the list.



Parameters

-s (Optional) Displays the storage enclosure ID. You cannot use the **-l** and the **-s** parameters together.

-l (Optional) Displays default output, plus the storage enclosure feature code and engineering change level. You cannot use the **-l** and the **-s** parameters together.

-state *normal | not normal*

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

Note: You cannot specify an ID range.

If you specify the dash (-), this parameter information is automatically supplied.

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 **lsstgenc1** command using the **-l** parameter. There are some differences in the returned values for a 2107 versus a 1750.

Invoking the **lsstgenc1** command

```
dscli>lsstgenc1 -l IBM.2107-75FA120
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI  
Version: 5.0.0.0 DS: IBM.2107-75FA120
```

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	RPM	State
0 x 0	16	16	145	10000	normal
0 x 1	16	16	145	10000	normal

Report field definitions

ID Specifies 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 Specifies the frame number within the designated storage unit that contains this storage enclosure.

Enclnum

(2107) Identifies a storage enclosure within a storage unit frame. The values that are displayed for the enclosure are S11 - S18 or S21 - S28. These values are interpreted as follows:

- Enclosures S1x are located in front of the frame.
- Enclosures S2x are located in the rear of the frame.
- Enclosure Sx1 is located at the top of the frame.
- Enclosure Sx8 is located at the bottom of the frame.

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

FC Specifies the feature code that was used to order this storage enclosure.

Interfaces

Specifies 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 b*, (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
- *b* = indicates that the value is a binary string

The values for DS6000 represent the following:

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

Specifies the FC-AL interface base address assigned to this storage enclosure for DDM access.

Storslot

Specifies the number of slots for storage devices in this storage enclosure.

Stordev

Specifies the number of storage devices that are installed in this storage enclosure.

Cap Specifies 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 Specifies 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 Specifies the condition of the storage enclosure. The condition of the enclosure is either *normal* or one of the conditions that falls under the category *not normal*. The following values can be displayed:

normal

Indicates that the storage enclosure is not in any failure or transient condition.

offline

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

failed Indicates that the storage enclosure is broken and ready to be removed without impacting the system.

resuming

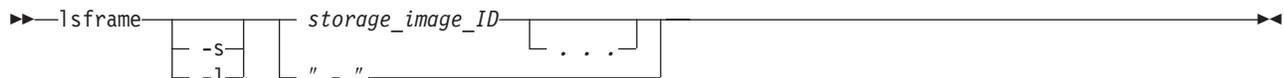
Indicates that the storage enclosure is in the process of coming online.

quiescing

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

lsframe

The **lsframe** command displays a list of frame enclosures for a storage image.



Parameters

-s Displays the rack enclosure ID. You cannot use the **-l** and the **-s** parameters together.

-l Displays default output and the frame enclosure location. You cannot use the **-l** and the **-s** parameters together.

storage_image_ID . . . | -

Displays frame enclosure information for the specified storage image IDs. A storage image ID includes manufacturer, type, and serial number. You must separate multiple IDs with spaces.

Note: ID ranges cannot be specified.

If you specify the dash (-), this parameter information is automatically supplied.

Example

Invoking the lsframe command

```
dsccli>lsframe -l IBM.2107-75FA120
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

ID	frm#	Frame	loc
IBM.2107-75FA120/R1	R1	1	U2107-75FA120
IBM.2107-75FA120/R2	R2	1	U2107-75FA120

Report field definitions

You can use the following information to help you understand the output that is generated from this command:

ID Identifies the unique identifier of the frame enclosure.

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

frameID 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.ppssss*.

Physical resource information commands

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

Use the following commands to view information about the physical resources in your storage complex:

- **lsddm**
- **lshba**
- **lsda**

The **lsddm** command displays a report that lists the disk drive modules and status information for each disk drive module in the list.

The **lshba** command displays a report that lists the storage image host adapters and status information for each host adapter in the list.

The **lsda** command displays a report that lists the device adapters (DA) for each storage image and the status of each device adapter in the list.

lsda

The **lsda** command displays a list of device adapters (DA) for each storage image. You can use this command to look at the status of each device adapter in the list.

```
▶▶—lsda [ -s ] [ -enc enclosure_ID ] [ -server server_ID ] [ -dapair dapair_ID ]
```


FC	Server	DA pair	Interfaces
1234	00	IBM.2107-75FA120/11	111,0112,0113,0114
1234	01	IBM.2107-75FA120/12	111,0112,0113,0114
1234	00	IBM.2107-75FA120/11	111,0112,0113,0114
1234	01	IBM.2107-75FA120/12	111,0112,0113,0114

Report field definitions

The following supplies information that you can use to help you understand the output that is generated from this command.

ID Specifies the unique identifier of the device adapter.

State Specifies the current availability state of the specified device adapter.

Loc Specifies the I/O enclosure and the device adapter location.

The I/O enclosure location format is *Uttt.mmm.ppssss*.

The device adapter location format is *Pn-Cn* where *Pn* indicates the Planner number (1) and *Cn* indicates the card number (1 - 6).

FC Specifies the feature code that is used to order the specified device adapter.

Server Specifies the server or device adapter group to which the device adapter is assigned.

DA pair

Specifies 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.

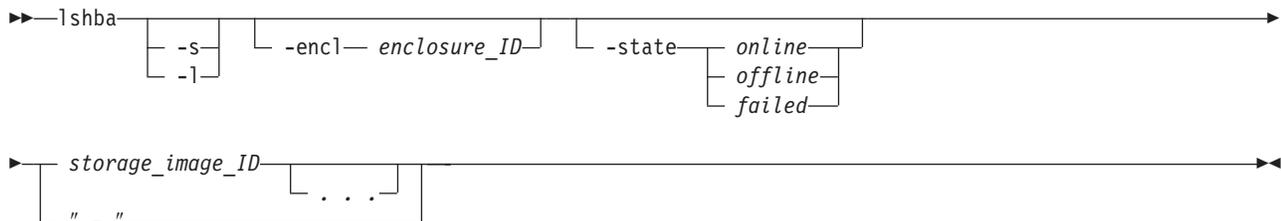
An even numbered device adapter pair ID indicates the first device adapter pair in an I/O enclosure pair. An odd numbered device adapter pair ID indicates the second device adapter pair in an IO enclosure pair.

Interfaces

Specifies the four interface IDs that are associated with the FC-AL ports.

Ishba

The **Ishba** command displays a list of storage image host adapters and status information for each host adapter in the list.



Parameters

- s (Optional) Displays host adapter IDs. You cannot use the -l and the -s parameters together.
- l (Optional) Displays the default output plus the host adapter feature code and interface IDs. You cannot use the -l and the -s parameters together.
- encl *enclosure_ID*
(Optional) Displays host adapters that are associated with a common processor complex or I/O enclosure ID.
- state *online | offline | failed*
(Optional) Displays host adapters that are in a specified state.
- storage_image_ID . . . | -*
(Optional) Displays host adapter information for the specified storage image IDs. A storage image ID includes manufacturer, model, and serial number. You must separate multiple IDs with spaces.

Note: ID ranges cannot be specified.

If you specify the dash (-), this parameter information is automatically supplied.

Example

Invoking the lshba command

```
dsccli>lshba -l IBM.2107-75FA120
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

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	Interfaces
1234	0x0121,0x0121,0x0123,0x0124
1234	0x0131,0x0131,0x0133,0x0134
1234	0x0221,0x0221,0x0223,0x0224
1234	0x0231,0x0231,0x0233,0x0234

Report field definitions

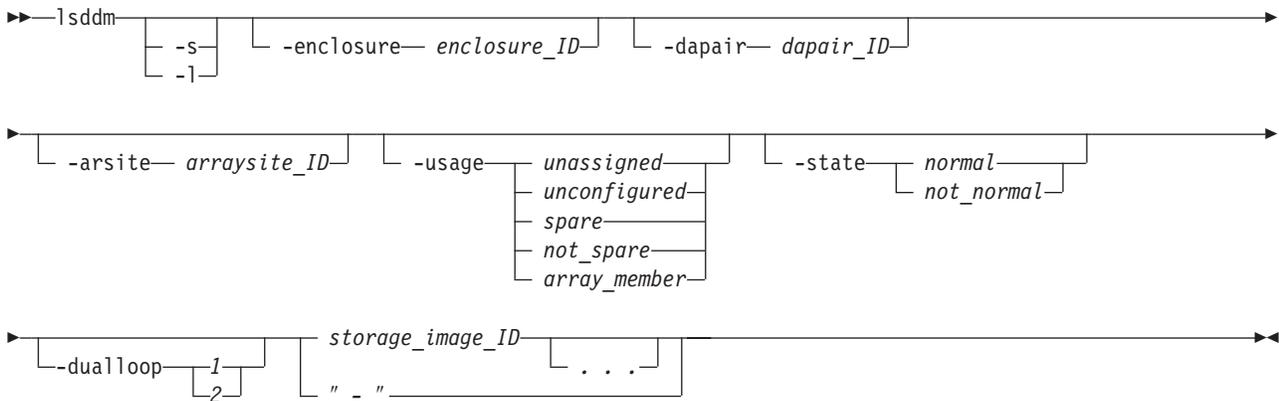
You can use the following information to help you understand the output that is generated from this command:

ID Specifies the unique identifier of the host adapter.

- State** Specifies the current availability state of the specified host adapter.
- Loc** Specifies the I/O enclosure and the host adapter location.
The I/O enclosure location format is *Uttt.mmm.ppssss*.
The host adapter location format is *Pn-Cn* where *Pn* indicates the Planner number (1) and *Cn* indicates the card number (1 - 6).
- FC** Specifies the feature code that is used to order the specified host adapter.
- Interfaces**
Specifies the four interface IDs that are associated with the I/O ports on the host adapter.

lsddm

The **lsddm** command displays a list of disk drive modules (DDMs) and status information for each DDM in the list.



Parameters

- s** (Optional) Displays the DDM IDs. You cannot use the **-s** and **-l** parameters together.
- l** (Optional) Displays the default output. You cannot use the **-s** and **-l** parameters together.
- enclosure enclosure_ID**
(Optional) 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) 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) 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) Displays DDMs that are associated with a specified usage.

-state *normal | not_normal*
 (Optional) Displays DDMs that are associated with a specified state.

-dualloop *1 | 2*
 (Optional) Specifies that DDMs that are associated with a specified dual loop be displayed.

storage_image_ID . . . | -
 (Required) Displays DDM information for the specified storage image IDs. A storage image ID includes manufacturer, type, and serial number. You can specify multiple IDs and they must be separated with a space between each ID.

Note: You cannot specify ID ranges.

If you use the dash (-), this parameter information is automatically supplied.

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 **lsddm** command using the **-I** parameter. A separate example is not shown for the 1750 because the information is the same for both. The only difference is the machine type designation, 2107 versus 1750. Also, the table shows only the first 2 DDMs associated with the specified storage unit.

Invoking the lsddm command

```
dscli>lsddm -I IBM.2107-75FA120
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

ID	Model	loc	Firmware level	DA pair	Dualloop	dkcap (10^9B)
IBM.2107-D21-75FA120/R1-S11-P1-D1	S0BE146	U2107.921.75FA120	3603	11	1	145
IBM.2107-D21-75FA120/R1-S11-P1-D2	S0BE146	U2107.921.75FA120	3603	11	1	145

Diskrpm	Dkinf	Dkrate	Dkuse	Arsite	Position	State	Diskclass
1000	FCAL	2	array member	S123	1	normal	ENT

Diskrpm	Dkinf	Dkrate	Dkuse	Arsite	Position	State	Diskclass
1000	FCAL	2	array member	S123	2	normal	ENT

Report field definitions

ID Specifies the system-assigned unique identifier for the DDM.

Model Specifies 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 GB.

loc Specifies 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

Specifies the level of firmware that is installed on the specified DDM.

DA pair

Specifies the device adapter pair ID. DA pairs are located 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

Specifies the dual loop that the specified DDM resides on. The value is either 1 or 2.

dkcap (10^9B)

Specifies the DDM raw capacity in gigabytes.

diskrpm

Specifies the DDM rpm. One of the following values is displayed:

- 10000
- 15000

dkinf Specifies the DDM interface type. One of the following values are displayed:

- FC-AL
- FC-NL

dkrate Specifies the DDM interface rate.

dkuse Specifies the DDM usage in an array site. One of the following values are displayed:

- unassigned
- unconfigured
- spare required
- spare not required
- array member

arsite Specifies the array site ID.

Position

Specifies the DDM position in an array site configuration of DDMs.

- State** Specifies the current DDM state. One of the following values are 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 has been identified.
 - Verifying**
The storage device is made accessible to the device adapter, its 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 has failed and an immediate repair action is required.
 - Failed - Deferred Service**
The storage device has failed and a repair action is not immediately required.
 - Removed**
The storage device is removed from the system and removal has been 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

Specifies the disk class as either high speed fibre channel disk drives or near-line disk drives. The value displayed is one of the following:

- ENT = Specifies enterprise and represents high speed fibre channel disk drives
- NL = Specifies near-line and represents ATA (FATA) disk drives

Storage complex configuration commands

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

Use the following commands to configure a storage complex and show storage complex information:

- **chsp**
- **setvpn**
- **showsp**

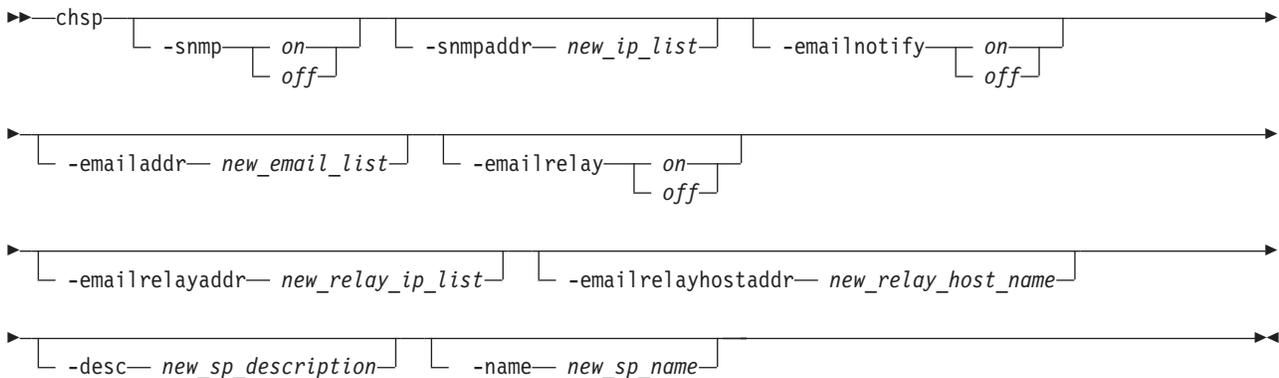
You can use the **chsp** command to modify a storage complex for items such as notification of the Simple Network Management Protocol (SNMP) traps and e-mail problem notification lists in a storage complex.

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.

The **showsp** command is used to generate 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 e-mail problem notification lists in a storage complex.



Parameters

Note: The Storage Manager server domain is a single storage complex. The storage complex object can only be created or deleted by service personnel.

-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.

-emailnotify *on | off*

(Optional) Specifies whether e-mail notifications are sent to the designated e-mail addresses when a problem occurs on a storage complex.

When you enable the **-emailnotify** and **-emailrelay** parameters, the e-mail notification is sent directly to the IP address that is associated with the **-emailrelayhost** *hostname* parameter. However, if the e-mail relay host name is not specified, the e-mail notification is sent directly to the e-mail relay address.

When e-mail notify is enabled and e-mail relay is disabled, the e-mail notification is sent directly to the specified e-mail address.

-emailaddr *new_email_list*

(Optional) Specifies the e-mail IP address or addresses where problem notifications are sent. You can specify two IP addresses.

-emailrelay *on | off*

(Optional) When you enable the **-emailrelay** and **-emailnotify** parameters, the e-mail notification is sent directly to the IP address that is associated with the **-emailrelayhost** *hostname* parameter. However, if the e-mail relay host name is not specified, the e-mail notification is sent directly to the e-mail relay address.

When e-mail notify is enabled and e-mail relay is disabled, the e-mail notification is sent directly to the specified e-mail address.

-emailrelayaddr *new_relay_ip_list*

(Optional) Specifies the e-mail relay IP address list. You can specify two IP addresses.

-emailrelayhostaddr *new_relay_host_name*

(Optional) Specifies the new e-mail relay host system name.

-desc *new_sp_description*

(Optional) Specifies your description of the storage complex. This description is limited to 256 byte or 128 double-byte characters.

-name *new_sp_name*

(Optional) Specifies the name that you designate for the storage complex. This name is limited to 32 byte or 16 double-byte characters.

Example

Invoking the chsp command

```
dscli>chsp -desc "my storage complex"
```

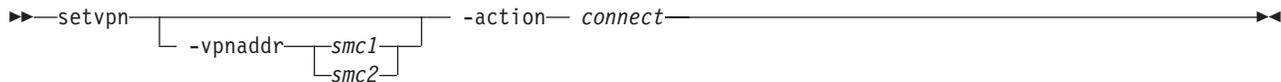
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
Storage-complex IBM.2107-75FA120 successfully modified.
```

setvpn

The **setvpn** command starts an outbound virtual private network connection.



Parameters

Notes:

1. 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.
2. It can take anywhere from 2 to 10 minutes for the secure connection to be established and visible from the RS3/RS4 server.
3. The secure connection is ended automatically when the terminal emulation session is ended.
4. 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.

-vpnaddr *smc1* | *smc2*

(Optional) Specifies the VPN Server Machine. In addition you can specify where the outbound VPN is to be started from by designating the following values:

smc1 Identifies the management console (SMC) you designate on the command line where the outbound VPN is to started from. Or you can leave this as *smc1*. The console that you have specified in your profile for *hmc1* is used unless you specify a console that is not designated in your profile to start your DS CLI session. 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 you designate on the command line where the outbound VPN is started from. You can also leave this as *smc2*. The console that you specified in your profile for *hmc2* is used unless you specify a console that is not designated in your profile to start your DS CLI session. In this case, the console that you specify to start your session is the one where the connection is made.

-action *connect*

(Required) Specifies that the secure VPN connection be started.

Example

Invoking the setvpn command

```
dscli>setvpn -vpnaddr smc1 -action connect
```

The resulting output

```
Sun July 10 02:23:49 PST 2004 DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120  
Secure connection started.
```

Example scenario

The following list provides a description of what the prerequisites are for using VPN:

- The customer must provide connectivity for the remote service using Internet VPN, or modem VPN.
- The customer must specify one of the following remote service methods:
 - Unattended (preferred): An IBM support person can initiate a remote service connection. This requires modem connectivity.
 - Attended: Only a person with physical access to the DS8000 HMC or the Storage Manager GUI can initiate a remote service connection.

Note: Although a modem is only required for the unattended remote service, the modem can serve as a backup for call home if Internet VPN connectivity does not work properly.

During the installation of the DS8000 product, the hardware management console (HMC) sends a certificate (signed public key) to IBM that is used for server authentication and for SSL encryption of applications using VPN (Internet and modem) connectivity.

After the prerequisites are completed, you are ready to use Internet VPN. The **setvpn** command is used to start the session and a secure connection is made. In addition, the IBM VPN server does additional authentication to allow traffic to certain IBM servers only, for call home and remote service.

Note: Only IBM support personnel with special access rights can use the VPN connection.

showsp

The **showsp** command (for DS8000) displays detailed properties of a storage complex. Detailed properties include the names, descriptions, and customer account names for the storage complex.

▶▶—showsp—▶▶

Parameters

There are no parameters for this command.

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

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
```

Name	Desc	Acct	SNMP	SNMPAddr
My_storage-complex	Production storage-complex	ABC Company	Enabled	9.xxx.14.245

eMailnotify	eMailaddr	eMailrelay	eMailrelay-addr	eMailrelay-hostaddr
Enabled	email1@ibmds8000.com, email2@ibmds1.com	Disabled	9.xxx.14.45	relay_host

Report field definitions

Name Specifies the name of the storage complex.

Desc Specifies the description of the storage complex.

Acct Specifies the customer account name for the storage complex.

SNMP

Specifies whether SNMP traps that are generated by the storage complex are sent to the IP address that is specified by the **chsp** command. This column displays the values: Enabled or Disabled.

SNMPAddr

Specifies 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

Specifies whether e-mail notifications are enabled or disabled.

When the **chsp** command enables the **-emailnotify** and **-emailrelay** parameters, the e-mail notification is directed to the IP address that is associated with the **-emailrelayhost hostname** parameter. However, if the e-mail relay host name is not specified, the e-mail notification is directed to the e-mail relay address.

When e-mail notify is enabled and e-mail relay is disabled, the e-mail notification is sent directly to the specified e-mail address.

eMailaddr

Specifies one or more e-mail addresses to which notification is sent if service is required when e-mail is enabled. A null (-) is displayed if e-mail is not enabled or if an e-mail address is not available.

eMailrelayaddr

Specifies one or more e-mail relay IP addresses. These IP addresses represent the addresses through which notification is relayed if service is required when e-mail is enabled.

eMailrelayhostaddr

Specifies the e-mail relay host system name.

Storage unit configuration commands

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

Use the following commands to configure a storage unit and show storage unit information:

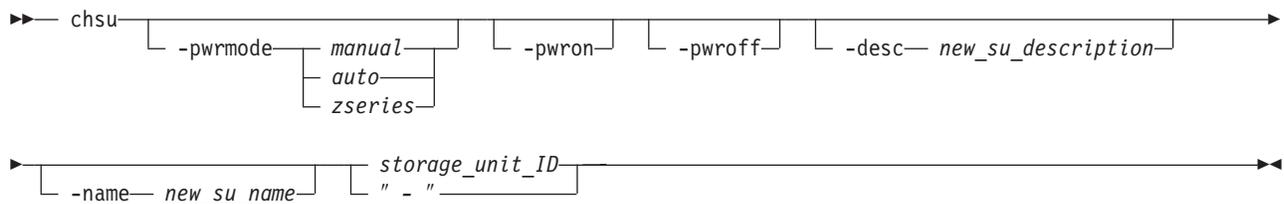
- **chsu**
- **lssu**
- **showsu**

The **chsu** command changes the description and name you have associated with a specified storage unit. You can also change the remote power control mode on the storage unit.

The **lssu** and **showsu** commands generate reports that allow you to view details about your storage units.

chsu

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



Parameters

-pwrmode *manual | auto | zseries*

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

manual

Specifies 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 - DS8000 only) Turns on power to the storage unit. This parameter is valid if the control mode is set to manual and the switch is set to remote.

-pwoff

(Optional - DS8000 only) Turns off power to the storage unit. This parameter is valid if the control mode is set to manual and the switch is set to remote.

-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.

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 specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the chsu command

```
dscli>chsu -pwrmode manual IBM.2107-75FA120
```

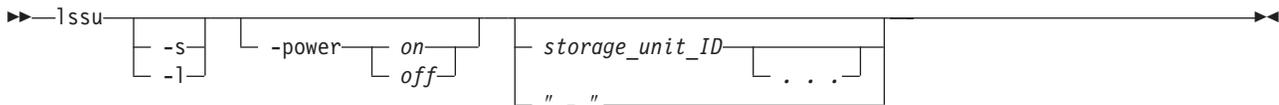
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
Storage unit IBM.2107-75FA120 successfully modified.
```

Issu

The **Issu** 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.



Parameters

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

-l (Optional) Displays default output and remote mode with scheduled power-on and power-off times, and the storage unit description. You cannot use the **-l** 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, type, and serial number. You must separate multiple IDs with a space between each ID.

Note: You cannot specify ID ranges.

If you specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for the 2107 and 1750. The only difference is the value input for the *storage_unit_ID* parameter.

Invoking the lssu command

```
dscli>lssu -l
```

The resulting output

Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

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	Production
SU 3	IBM.2107-75FA130	921	3007ACF 3078999 E0	On	Local	Backup

Report field definitions

Name Specifies the user-defined name for each storage unit found within the storage complex. This value is null (-) if you have not assigned a name to a storage unit.

ID Specifies 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.

Model

Identifies the model number of the storage unit.

WWNN

Specifies the World Wide Node Name for the listed storage unit. This value is null (-) 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.

Off

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 Scheduled

Indicates that the SMC local/remote switch is set to remote and that the power control mode is set to scheduled 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 SMC Scheduled/Auto

Indicates that the SMC local/remote switch is set to remote and that the power control mode is set to auto/scheduled 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 Specifies the description that you assigned the storage unit. This value is displayed as a null (-) if no description has been assigned.

showsu

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

```
▶▶—showsu— [ storage_unit_ID ]—————▶▶
```

Parameters

storage_unit_ID

(Required) Specifies the storage unit ID. A storage unit ID consists of manufacturer, machine type, and serial number.

If you specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the showsu command

```
dscli>showsu DS: IBM.2107-75FA120
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

Name	Desc	ID	Model	WWNN	Config
My Storage Unit	This is my DS Storage Unit	IBM.2107-75FA120	921	3007ACF3012399E0	One I/O interface

Pw State	Pw Mode	Reqpm	Processor Memory	MTS
On	Local	Remote SMC manual	1 GB	IBM.2424-75FA120

Report field definitions

Name Specifies the name that you assigned for the designated storage unit. This value is null (-) if you have not assigned a name to a storage unit.

Desc Specifies the description that you assigned for the designated storage unit. This value is displayed as a null (-) if no description has been assigned.

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

Model Identifies the model number of the designated storage unit.

WWNN

Specifies the World Wide Node Name for the listed storage unit. This value is null (-) if the WWNN is not known.

Config

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

Undefined

Indicates that a configuration upgrade is in progress that causes the configuration option to change.

One I/O interface

Indicates that there is one dedicated I/O interface between the I/O enclosure pairs and storage enclosures.

Two I/O interfaces

Indicates that there are two dedicated I/O interfaces between the I/O enclosure pairs and storage enclosures.

Four I/O interfaces

Indicates that there are four dedicated I/O interfaces between the I/O enclosure pairs and storage enclosures.

Six I/O interfaces

Indicates that there are six dedicated I/O interfaces between the I/O enclosure pairs and storage enclosures.

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.

Off

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 Scheduled
- Remote SMC Auto
- Remote SMC Scheduled/Auto
- Remote zSeries Power Control

Note: The default value is remote SMC Manual mode.

Processor Memory

Specifies the amount in gigabytes of processor memory configured on the storage unit.

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 ordered

before release 2.4. After release 2.4, the order type varies according the warranty periods associated with the storage unit.

Storage image configuration commands

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

Use the following commands to configure a storage image and show storage image information:

- **chsi**
- **lsserver**
- **lssi**
- **shows**

The **chsi** command primarily enables or disables the ESSNet user interface that invokes Copy Services operations for the storage image, changes the description and name that you have assigned to the storage image, or changes an iSystem serial number.

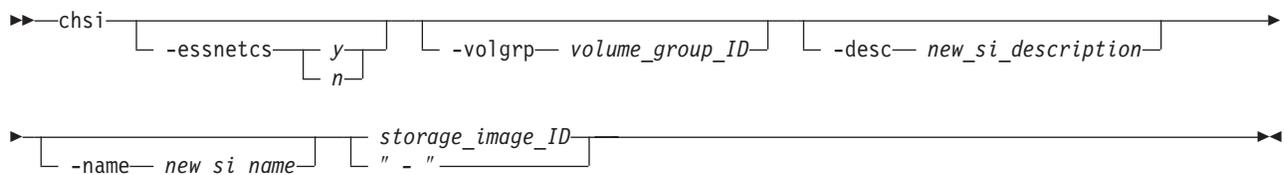
The **lsserver** command displays all the servers in a storage complex or a list of specified servers. The displayed list also provides the status information for each server including the LIC version number, operating system version number, and bundle version.

The **lssi** command with the **lsavailpprcport** and the **mkpprcpath** commands 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 **lssi** command.

The **shows** command 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)

When the **-essnetcs y** parameter is used, the **-volgrp** parameter specifies the ESSNet Copy Services type volume group that contains the logical volumes that are eligible for control by Copy Services operations. If the **-essnetcs y** parameter and the volume group ID are not specified, then all logical volumes are eligible for control by Copy Services operations.

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 byte 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.

storage_image_ID | -

(Required) Specifies the storage image ID, which consists of the values for manufacturer, machine type, and serial number.

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the chsi command

```
dscli>chsi -essnetcs n IBM.2107-75FA120
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

```
Storage image IBM.2107-75FA120 successfully modified.
```

lsserver

The **lsserver** command displays all servers in a storage complex or a list of specified servers and it also displays the status information for each server in the list.



Parameters

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

-l (Optional) Displays the default output and the state mode of the servers. You cannot use the **-l** and the **-s** parameters together.

Server_ID . . . | -

(Optional) Displays the server information for the specified server IDs. This parameter accepts a fully qualified server ID, which includes the storage image ID or a shortened version without the storage image ID. The shortened version is a two-digit decimal number with no leading zeros.

Example: IBM.2107-13AAV3A/0

Example: IBM.2107-13AAV3A/1

Note: If you choose to use this parameter with a fully qualified storage image ID, ensure that you properly specify the server ID with the /x or /xx server designation. A properly coded **lsserver** command that designates the server ID looks like the following:

```
dscli>lsserver -l IBM.2107-75FA120/0
```

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.

If you specify the dash (-), this parameter information is automatically supplied.

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 **-l** parameter. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750. The only difference is the value you input when you use the *Server_ID* parameter.

Invoking the lsserver command

```
dscli>lsserver -l
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM  
DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

ID	Image ID	Image Name	Power Control SFI	State	LIC Version	OS Version	Bundle Version
IBM.2107-75FA120/00	1	SF75ALK COESS01	0	Online	5.0.7.155	2005.9. 1.163506	6.0.500.8

Report field definitions

ID Specifies the unique identifier of the server. This value includes the storage image ID and the server ID.

Image ID

Specifies the image ID for the designated storage server.

Image Name

Specifies the image name for the designated storage server.

Power Control SFI

Specifies the storage server power control SFI.

LIC Version

Specifies the LIC version for the designated storage server.

State Specifies the current state of the designated server.

OS Version

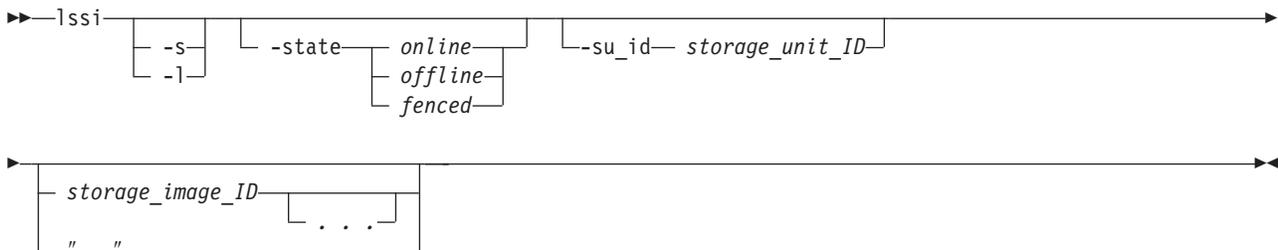
Specifies the operating system version for the designated server.

Bundle Version

Specifies the bundle version for the designated storage server.

Issi

The **Issi** command displays a list of storage images in a storage complex. You can use this command to look at the status of each storage image in the list. The storage image worldwide node name (WWNN) is displayed when this command is used. You must use the storage image WWNN with the **lsavailpprcport** and **mkpprcpath** commands.



Parameters

- s** (Optional) Displays storage image IDs. You cannot use the **-l** and the **-s** parameters together.
- l** (Optional) Displays the default output, ESSNet, volume group, and storage image description. You cannot use the **-l** and the **-s** parameters together.
- state** *online* | *offline* | *fenced*
(Optional) Displays only the storage images in the specified state.
- su_id** *storage_unit_ID* . . .
(Optional) Displays the storage images that are associated with the specified storage 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.

If you specify the dash (-), this parameter information is automatically supplied.

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 represents the headers that are displayed on the output report that is associated with the **Issi** command using the **-l** parameter. A separate

example is not shown for the 1750 because the information is the same for the 2107 and 1750. 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

Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

Name	ID	Storage Unit	Model	WWNN
DS 1	IBM.2107-75FA120	IBM.2107-75FA120	921	3007ACF3012399E0

State	ESSNet	Volume Group	Desc
Online	Enabled	-	This is my DS storage Image

Report field definitions

Name Specifies the name that you assigned to 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.

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

WWNN

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

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

Online

Indicates that the storage unit is capable of processing all functions.

Offline

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

Resuming

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

Quiescing

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

Quiesce Exception

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

Forced Quiescing

Indicates that the storage unit is in the process of performing a force offline operation.

Fenced

Indicates that the storage unit has failed and is offline.

ESSNet

Specifies that the storage-complex ESSNet user interface can invoke 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 null (-) 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 null (-) value in this field), all logical volumes are eligible to be controlled by Copy Services operations that are initiated through the ESSNet.

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

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 specify the dash (-), this parameter information is automatically supplied.

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 represents the headers that are displayed on the output report that is associated with the **showsi** command. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750. The report displays the iSeries serial number for a 1750 machine type and an MTS column for a 2107.

Invoking the showsi command

```
dsccli>showsi
IBM.2107-75FA120
```

The resulting output

Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

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	MTS
Online	Enabled	-	-	8 GB	128 GB	1 GB	IBM.2421-75FA120

Report field definitions

Name Specifies the name that you assigned to the storage unit.

Desc Specifies the value that you have 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.

Model 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-xxxx*.

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

Online

Indicates that the storage unit is capable of processing all functions.

Offline

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

Resuming

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

Quiescing

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

Quiesce Exception

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

Forced Quiescing

Indicates that the storage unit is in the process of performing a force offline operation.

Fenced

Indicates that the storage unit has failed and is offline.

ESSNet

Specifies that the storage-complex ESSNet user interface can invoke 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 null (-) 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 null (-) value in this field), all logical volumes are eligible to be controlled by Copy Services operations that are initiated through the ESSNet.

OS400Serial (1750 only)

Specifies null (-) for a 2107 model type and the iSeries serial number for a 1750 model type.

The serial number consists of 3 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 directed to each LUN.

NVS Memory

Specifies the amount in gigabytes of nonvolatile storage (NVS) memory configured on the storage unit.

Cache Memory

Specifies the amount in gigabytes of cache memory configured on the storage unit.

Processor Memory

Specifies the amount in gigabytes of processor memory configured on the storage unit.

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 the warranty periods that are associated with the storage unit.

I/O port and host connection configuration commands

This section contains storage image I/O port and host connection configuration commands.

Use the following commands to configure your storage image I/O ports and host connections and to display connection information.

Storage image I/O port commands

This section contains commands that are used to configure storage image I/O port information and to display storage image I/O port information.

Use the following commands to configure storage image I/O ports and to display storage image I/O port information:

- **lσιοport**
- **setιοport**
- **showιοport**

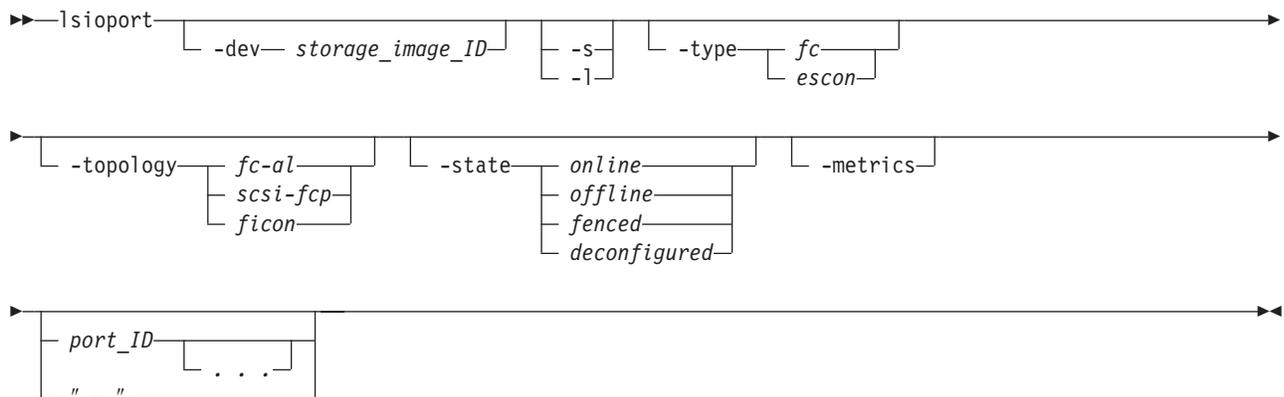
The **lσιοport** command displays a list of I/O ports on a specified storage image and optionally provides performance metrics for each I/O port that is listed.

The **setιοport** command configures one or more I/O ports for open systems or zSeries host system connections.

The **showιοport** command displays the properties of a specified I/O port. It optionally displays the performance metrics for the I/O port.

lσιοport

The **lσιοport** command 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.



Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which consists of manufacturer, type, and serial number.

Note: This parameter is required if you do not specify a fully qualified port ID and you do not specify a value for the *devid* variable in your profile file.

-s (Optional) Displays fully qualified port IDs. You cannot use the **-l** and **-s** parameters together.

-l (Optional) Displays default output plus the maximum login limit and the number of current logins. You cannot use the **-l** and **-s** parameters together.

-type *fc* | *escon*

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

-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 unit 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:

- *EE* is an I/O port enclosure number in the range of 00 - 17 (2107 machine types).
- *A* is the adapter number and is specified as 1, 2, 4, or 5 (2107 machine types).
- *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.

If you specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the lsioport command

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

The resulting output

Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

ID	WWPN	State	Type	Topo	Portgrp	Speed
IBM.2107-75FA120/I0111	307BCF30A3299E0	Online	Fibre Channel-LW	SCSI-FCP	0	1Gb/s

Report field definitions (without the -metrics parameter)

ID Specifies the fully qualified port ID, which includes the storage image ID.

WWPN

Specifies the fibre channel worldwide port number. If the port type is not fibre channel, the value displayed is null (-).

State Specifies the current I/O port status. One of the following values can be displayed:

Online

Indicates that the storage unit is capable of processing all functions (default).

Offline

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

Resuming

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

Quiescing

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

Fenced

Indicates that the storage unit has failed and is offline.

Deconfigured

Indicates that the I/O port is in the process of being deleted.

Type Specifies the port type. The following values can be displayed:

- Ethernet - Copper
- Fibre Channel-SW - (SW stands for short wave)
- Fibre Channel - LW (4 KM)
- ESCON - LED (2105 only)

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

- FC-AL
- SCSI-FCP
- FICON
- Null (-) This value is displayed when the port type is not fibre channel.

Portgrp

Specifies 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 Specifies the I/O port interface speed. The following values can be displayed:

- ESCON ports = 200Mb/s
- FCP ports = 1 Gb/s, 2 Gb/s, 4 Gb/s
- FICON ports = 1 Gb/s, 2 Gb/s, 4 Gb/s

Report field definitions (with -metrics parameter)

When you use the **-metrics** parameter and do not specify a port type, two reports are displayed. 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. A report is not displayed for a port type that has no ports that are enabled.

FICON/ESCON ports: Each of the following headers and value types are displayed:

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 format that is used to report this value.

byteread
Specifies the number of bytes that are read in 128 KB increments.

bytewrit
Specifies the number of bytes that are written in 128 KB increments.

Reads
Specifies a value that is based on extended count-key-data (ECKD) data received operations.

Writes Specifies a value that is based on ECKD data transferred operations.

Timeread
Specifies 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.

Timewrite
Specifies 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.

SCSI-FCP ports: Each of the following headers and value types are displayed:

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 format used to report this value.

Bytewrit
Specifies a value for the remote mirror and copy data transferred operations in increments of 128 KB.

Bytoread
Specifies a value for the remote mirror and copy data received operations in increments of 128 KB.

Writes Specifies a value for the remote mirror and copy data transferred operations.

Reads Specifies a value for the remote mirror and copy data received operations.

Timewrite
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
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.

Bytoread
Specifies a value that is based on the SCSI data received operations. The displayed value is based on increments of 128 KB.

Reads

Specifies a value that is based on the SCSI data transferred operations.

Writes Specifies a value that is based on the SCSI data transferred operations.

Timeread

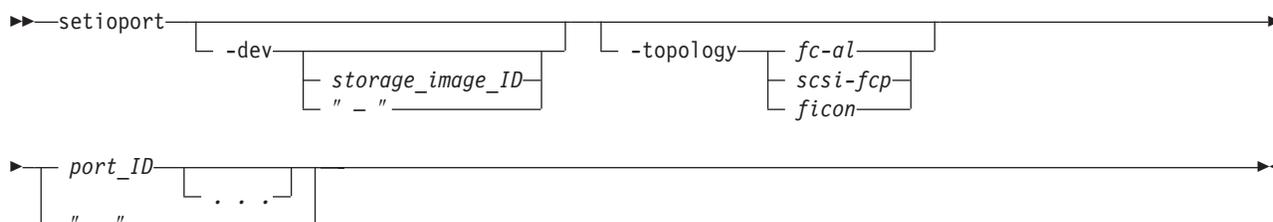
Specifies a value that is based on the SCSI data received (read-accumulated) time on a channel. The displayed value is based on increments of 16 milliseconds.

Timewrite

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.

setioport

The **setioport** command configures one or more I/O ports for open systems or zSeries host system connections. This command cannot be used for ESCON ports.



Parameters

-dev *storage_image_ID* | -

(Optional) Accepts a fully qualified storage image ID. The storage image ID consists of manufacturer, machine type, and serial number.

Note: This parameter is required if you do not specify a fully qualified port ID and you do not specify a value for the *devid* variable in your profile file.

-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.

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) Displays I/O ports matching the specified IDs. 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.

parameter accepts a fully qualified unique port ID, that is represented in the following format: manufacturer.machine type-serial number/portID.

For example, IBM.2107-75FA120/IO110

A port ID is prefixed with the letter *I* and consists of four hexadecimal characters in the format *EEAP*, where:

- *EE* is an I/O port enclosure number in the range of 00 - 17 (2107 machine types).
- *A* is the adapter number and is specified as 1, 2, 4, or 5 (2107 machine types).
- *P* is the port number (0 - 3).

If you specify the dash (-), this parameter information is automatically supplied.

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 **showioport** command. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the showioport command to show port information

```
dscli>showioport -dev IBM.2107-75FA120 IO112
```

The resulting output

Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS:IBM.2107-75FA120

ID	WWPN	State	Loc	Type
IBM.2107-75FA120/IO112	307ACF30A2399E0	Online	U2107-75FA123-I1-P1-C2-T1	Fibre channel-LW

Speed	Topo	Portgrp	unkSCSIlog
2 Gb/s	SCSI-FCP	0	-

Report field definitions (without the -metrics 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 null (-).

State Specifies the current state of the I/O port. One of the following values is displayed:

Online

Indicates that the storage unit is capable of processing all functions (default).

Offline

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

Resuming

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

Quiescing

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

Fenced

Indicates that the storage unit has failed and is offline.

Deconfigured

Indicates that the I/O port is in the process of 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*.

Type Specifies the port type. The following values can be displayed:

- Ethernet - Copper
- Fibre Channel-SW - (SW stands for short wave)
- Fibre Channel - LW 4KM
- ESCON - LED (2105 only)

Speed Specifies the I/O port interface speed. The following values can be displayed:

- ESCON ports = 200Mb/s
- FCP ports = 1 Gb/s, 2 Gb/s, 4 Gb/s
- FICON ports = 1 Gb/s, 2 Gb/s, 4 Gb/s

Topo Specifies the port topology. If the port type is not fibre channel, then the displayed value is null (-). One of the following values is displayed:

- FC-AL
- SCSI-FCP
- FICON
- - (null if not fibre channel)

Portgrp

Specifies an identifier that associates a subset of the I/O port objects that are operating in anonymous access mode.

unkSCSIlog

Specifies a list of unknown SCSI N-port WWPN identifiers that have attempted to login into this I/O port.

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

Specifies the number of bytes that are read in increments of 128 KB.

bytewrit

Specifies the number of bytes that are written in increments of 128 KB.

Reads

Specifies a value that is based on the extended count-key-data (ECKD) architecture data received operations.

Writes Specifies a value that is based on the ECKD architecture data transferred operations.

Timeread

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

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.

Bytewrit

Specifies a value for the remote mirror and copy data transferred operation in increments of 128 KB.

Byteread

Specifies a value for the remote mirror and copy data received operations in increments of 128 KB.

Writes Specifies a value for the remote mirror and copy data transferred operations.

Reads Specifies a value for the remote mirror and copy data received operations.

Timewrite

Specifies a value 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

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

Specifies a value that is based on the SCSI data received operations. The displayed value is based on increments of 128 KB.

Reads

Specifies a value that is based on the SCSI data received operations.

Writes Specifies a value that is based on the SCSI data transferred operations.

Timeread

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

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.

Host connection commands

This section contains commands that are used to configure host connections and to display host connection information.

Use the following commands to configure host connections and to display host connection information:

- **chhostconnect**
- **lshostconnect**
- **lshostvol**
- **lspportprof**
- **managehostconnect**
- **mkhostconnect**
- **rmhostconnect**
- **showhostconnect**
- **lshosttype**

The **chhostconnect** command 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.

The **lshostconnect** command displays a list of host connections for a storage image and the status information for each host connection in the list.

The **lshostvol** command displays the mapping of host device names or volume names to machine type 2105, 2107, and 1750 volume IDs.

The **lspportprof** command 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 particularly helpful to obtain the recommended values for the **mkhostconnect** command.

The **managehostconnect** command 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.

The **mkhostconnect** command configures the open systems hosts port attachments to fibre-channel ports that are configured for FC-AL or SCSI-FCP topology.

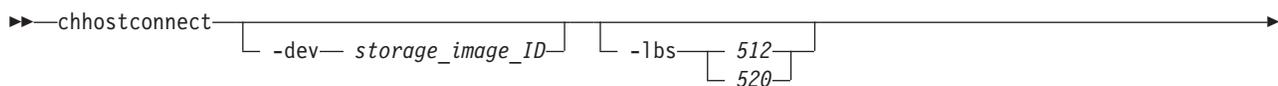
The **rmhostconnect** command removes a SCSI host port connection from a storage image.

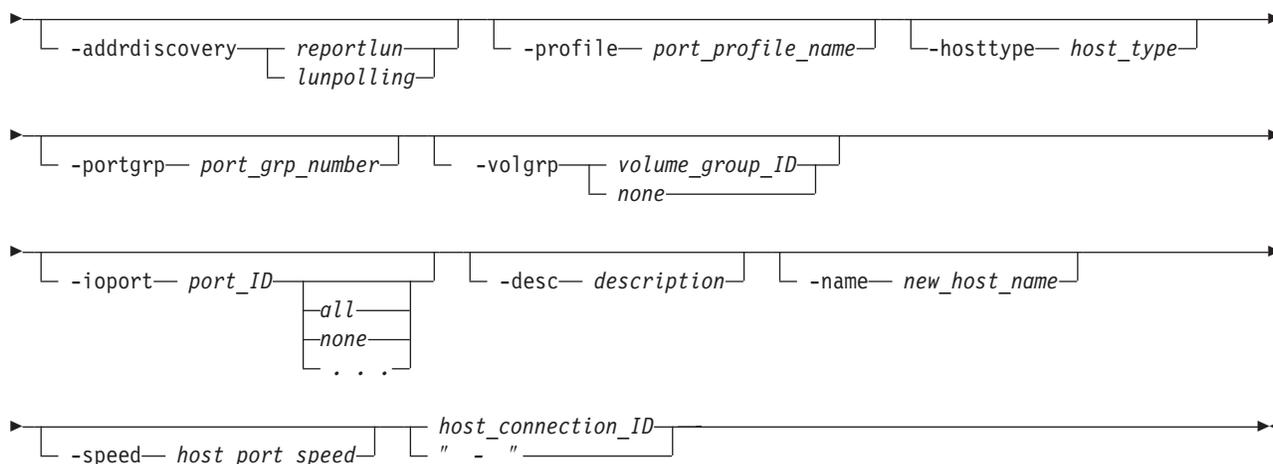
The **showhostconnect** command displays the detailed properties of a specified storage image host connection.

The **lshosttype** command 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.





Parameters

Notes:

1. The **chostconnect** command can be disruptive to host system I/O operations if the affected host port is logged into the target storage unit. You must ensure that the host port is offline to the host system before you process the **chostconnect** 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**
 - **-lbs**
3. If you are using the HP-UX operating system, see the volume restriction that is described under the **-addrdiscovery** parameter.

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which consists of the manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified host connection ID.

Example of a fully qualified storage image ID: IBM.2107-75FA120

-lbs *512 | 520*

(Optional) Specifies the logical block size that is used by the host system. The block size must be compatible with the volume group type and the volume type configurations that apply to the host port connection. The 520 logical block size is typically used by the IBM iSeries systems (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 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 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**
- **-lbs**

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*.

Examples of -dev parameter use

If you specify the **-dev** parameter, you can use the shortened version of the **-volgrp** parameter as follows:

```
dsccli>chhostconnect -dev IBM.2107-75FA120 V11 1
```

where *1* represents the required parameter, *host_connection_ID*.

If you do not specify the **-dev** parameter and you specify the **-volgrp** parameter, you must use the fully qualified version of the volume group ID with the **-volgrp** parameter specified as follows:

```
dscli>chhostconnect -volgrp IBM.2107-75FA120/V11 IBM.2107-75FA120/1
```

where *IBM.1750-68FA120/1* or *IBM.2107-75FA120/1* represents the required parameter, *host_connection_ID*

A host connection can use only one volume group per storage image. In other words, 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 all|none| . . .*

(Optional) Specifies all, none, 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.

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.

... (ellipsis)

Specifies that you can designate up to 128 ports for an open systems host attachment assignment. If you enter a list of I/O port IDs, access from the specified host connection to the specified volume group is allowed using only the designated list of port IDs.

Examples of **-dev** parameter use

If you specify the **-dev** parameter, you can use the shortened version of the **-ioport** parameter as follows:

```
dscli>chhostconnect -dev IBM.2107-75FA120 -ioport I0222 1
```

where *1* represents the required parameter, *host_connection_ID*.

If you do not specify the **-dev** parameter and you specify the **-ioport** parameter, you must use the fully qualified version of the port ID with the **-ioport** parameter specified as follows:

```
dscli>chhostconnect -ioport IBM.2107-75FA120/I0222 IBM.2107-75FA120/1
```

where *IBM.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:

- *EE* is an I/O port enclosure number in the range of 00 - 17 (2107 machine types).

- *A* is the adapter number and is specified as 1, 2, 4, or 5 (2107 machine types).
- *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.

-speed *host_port_speed*

(Optional) Specifies the host I/O port interface speed. This value defaults to *unknown* if it is not specified.

Note: If you want to explicitly specify a value of unknown, use 0 (zero) for the value.

You might want to specify the advertised interface speed which is generally one of the following values: 1Gb/s, 2Gb/s, 4Gb/s or 200Mb/s.

You can issue the **lshostconnect** command to view the value recognized within the system. You might find that the value specified on the report that is associated with the **lshostconnect** command is different than the value you have specified with this parameter. The system value overrides your input if there is a difference.

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.

Examples of -dev parameter use

If you specify the **-dev** parameter, you can use the shortened version of the **host_connection_ID** parameter as follows:

```
dscli>chhostconnect -dev IBM.2107-75FA120 -desc newdescription 1
```

where *1* represents the value for the *host_connection_ID*.

If you do not specify the **-dev** parameter and you specify the **host_connection_ID** parameter, you must use the fully qualified version of the host connection ID as follows:

```
dscli>chhostconnect -desc newdescription IBM.2107-75FA120/1
```

If you specify the dash (-), this parameter information is automatically supplied.

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 make a change to connection ID 0005 and type 000, the system makes the change to connection ID 0. Or, if you want to make a change to connection ID 0020 and type 002, the system makes the change to 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.

Example (2107)

Invoking the `chhostconnect` command

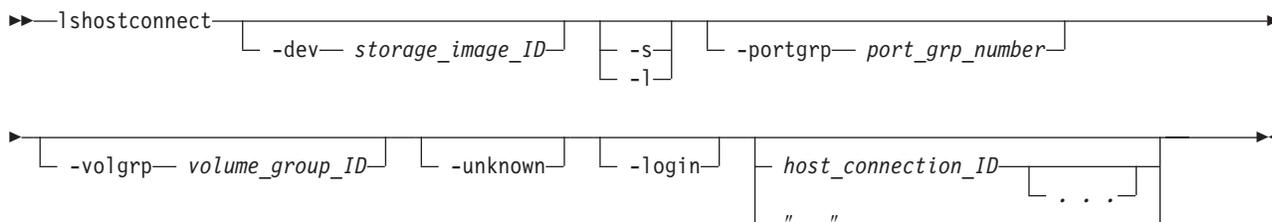
```
dscli>chhostconnect
-dev IBM.2107-75FA120
-name host_1_port_2 1
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
Host connection 1 successfully modified.
```

`lshostconnect`

The `lshostconnect` command displays a list of host connections for a storage image and the status information for each host connection in the list. You can also use this command to obtain a list of worldwide port numbers (WWPNs) from a system-detected-unknown host port. You can use these WWPNs to create a new 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. This parameter is required if you do not specify a fully qualified host connection ID and you do not specify a value for the *devid* variable in your profile file.
Example of a fully qualified storage image ID: IBM.2107-75FA120
- s** (Optional) Specifies the host connection IDs. You cannot use the **-l** and **-s** parameters together.
- l** (Optional) Specifies the default output and your description for each host connection in the list. You cannot use the **-l** and **-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 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. 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 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 those that you have created a host connection for, 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.

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.

If you specify the dash (-), this parameter information is automatically supplied.

Example of a fully qualified host connection ID: IBM.2107-75FA120/2

Example of a shortened version host connection 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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the lshostconnect command without the -unknown parameter

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

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

Name	ID	WWPN	HostType	LBS	Addrdiscovery
My host port 1	IBM.2107-75FA120/1	3007ACF30A2399E0	Unknown	512	reportLUN
My host port 2	IBM.2107-75FA120/2	3007ACF30A2399E1	Unknown	512	reportLUN
My host port 3	IBM.2107-75FA120/3	3007ACF30A2399E2	Unknown	512	reportLUN
My host port 4	IBM.2107-75FA120/4	3007ACF30A2399E3	Unknown	512	reportLUN

Profile	portgrp	volgrpID	achtopo	ESSIOport	Speed	Desc
IBM pSeries - AIX	0	V100	SCSI-FCP	I0111,I0121 I0211,I0221	1 GB/s (null (-) is reported for a 1750 model)	SCSI1
IBM pSeries - AIX	0	V100	SCSI-FCP	All	1 GB/s	SCSI2
IBM pSeries - pLinux	0	V100	SCSI-FCP	I0111,I0121 I0211,I0221	1 GB/s (null (-) is reported for a 1750 model)	SCSI3
IBM pSeries - pLinux	0	V100	SCSI-FCP	-	1 GB/s	SCSI4

Example of lshostconnect using the -unknown parameter

Invoking the lshostconnect command with the -unknown parameter

```
dsccli>lshostconnect
-dev IBM.2107-75FA120 -unknown
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

WWNN	WWPN	ESSIOport
3007ACF30A239900	3007ACF30A2399E0	I0111, I0121, I0211, I0221
3007ACF30A239900	3007ACF30A2399E1	I0121
3007ACF30A239900	3007ACF30A2399E2	I0111, I0121, I0211, I0221
3007ACF30A239900	3007ACF30A2399E3	I0111

Example of lshostconnect using the -login parameter

Invoking the lshostconnect command with the -login parameter

```
dsccli>lshostconnect
-dev IBM.2107-75FA120 -login
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

WWNN	WWPN	ESSIOport	LoginType	Name	ID
3007ACF30A239900	3007ACF30A2399E0	I0111	SCSI	MyHostA	1
3007ACF30A239900	3007ACF30A2399E1	I0111	SCSI	MyHostB	1
3007ACF30A239900	3007ACF30A2399E2	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

Specifies the worldwide port name (WWPN) for this host system port.

HostType

Specifies 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

Specifies 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 iSeries host system attachments.

Addrdiscovery

Specifies 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

Specifies the name of the host connection behavior profile.

Portgrp

Specifies 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

Specifies 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.

Achtopo

Specifies the topology of the attached unit. This value is used to select which storage facility I/O ports are compatible with the attachment to the specified storage unit. One of the following values is displayed:

- Unknown
- SCSI-FCP
- FC-AL

ESSIOport

Specifies the array of port IDs that the designated host port is logged into.

The port ID component is four hexadecimal characters in the format *EEAP*, where *EE* is a port enclosure number 00 - 17 for 2107, *A* is the adapter number 1, 2, 4, 5 for 2107, and *P* is the port number (0 - 3). The number is prefixed with the letter *I*.

Speed

Specifies the system identified host I/O port interface speed. This field displays a value of *unknown* if the interface speed was not specified. You might find that the value specified on the report is different than the value you specified with the **chhostconnect** command. The system value overrides your input if there is a difference.

Note: This value is not reported for a 1750 model type. A null (-) value is displayed.

Desc

Specifies the description that you defined for the SCSI host port. The description is limited to 256 byte or 128 double-byte characters.

Report field definitions when the -unknown parameter is used

WWNN

Specifies the worldwide node name for the designated host system.

WWPN

Specifies the worldwide port name for the designated host system port.

ESSIOport

Specifies the array of port IDs that the designated host port is logged into.

The port ID component is four hexadecimal characters in the format *EEAP*, where *EE* is a port enclosure number 00 - 17 for 2107, *A* is the adapter number 1,2,4,5 for 2107, and *P* is the port number (0 - 3). The number is prefixed with the letter *I*.

Report field definitions when the -login parameter is used

WWNN

Specifies the worldwide node name (WWNN) for this host system.

WWPN

Specifies the worldwide port name (WWPN) for this host system port.

ESSIOport

Specifies the port ID that the designated host port is logged into.

The port ID component is four hexadecimal characters in the format *EEAP*, where *EE* is a port enclosure number 00 - 17 for 2107, *A* is the adapter number 1,2,4,5 for 2107, and *P* is the port number (0 - 3). The number is prefixed with the letter *I*.

LoginType

Specifies the type of login such as SCSI.

Name

Specifies the name that you assigned to the host. If a name is not assigned, a null (-) 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.

lshostvol

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 MSCP paths.

- 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.

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the lshostvol command

```
dscli>lshostvol
```

The resulting output with SDD installed

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

Device/Volume 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

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

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	-
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 null (-).

lsportprof

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 specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the lsportprof command

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

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

Name	AddrDiscovery	LBS
IBM pSeries – AIX	ReportLUN	512

Name	AddrDiscovery	LBS
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

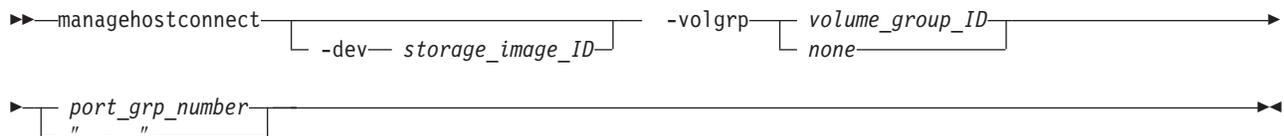
Specifies that host system LUN access is limited to a maximum of 64000 LUNs

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.



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 the manufacturer, type, and serial number.

-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 specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the managehostconnect command

```
dscli>managehostconnect -dev IBM.2107-75FA120 -volgrp 11 1
```

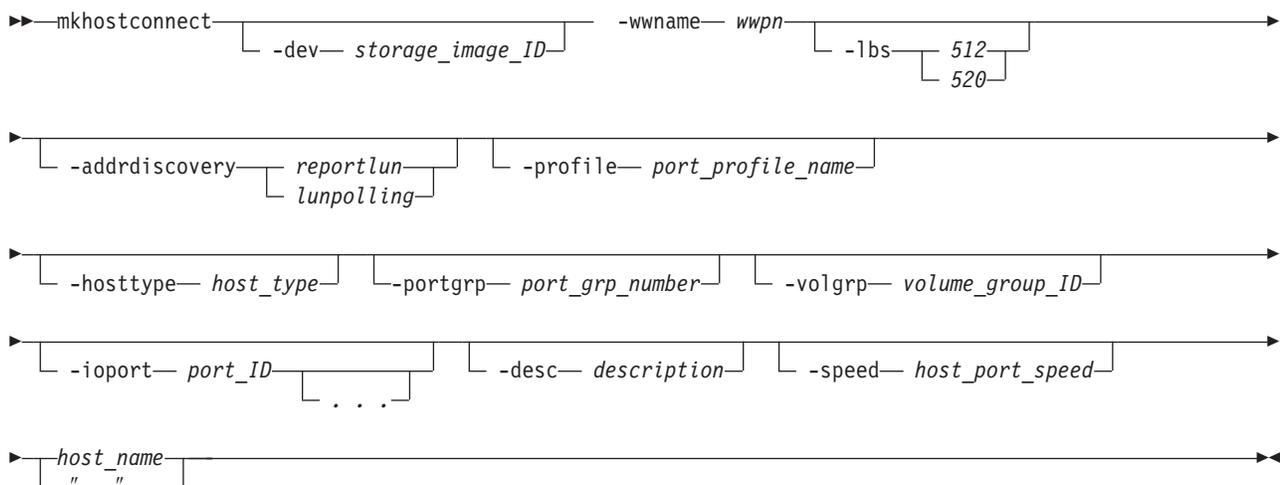
The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
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. Open systems hosts port attachments to fibre-channel ports are configured for identified access mode and SCSI protocol.



Parameters

Notes:

1. It is recommended 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**
- **-lbs**

2. If you are using the HP-UX operating system, see the volume restriction that is described under the **-addrdiscovery** parameter.

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes the manufacturer, type, and serial number.

-wwname *wwpn*

(Required) Specifies the worldwide port name (WWPN). The WWPN is a 16-character hexadecimal ID. The names are host attachment specific; for example, 12341234000A000F.

-lbs *512 | 520*

(Optional) Specifies the logical block size that is used by the specified host system, in bytes. The block size must be compatible with the volume group type and the volume type configurations that apply to the specified host port connection. The 520-byte size is typically used by IBM iSeries systems (OS/400®).

Notes:

1. Do not use the **-lbs** parameter if you specify the **-hosttype** parameter.
2. If you do not use the **-hosttype** parameter, use the **lspportprof** 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 **lspportprof** command to obtain a list of available profiles.

-hosttype *host_type*

(Optional) Specifies information about the following three parameters:

- **-profile**
- **-addrdiscovery**
- **-lbs**

Notes:

1. When the **-hosttype** parameter is specified, do not use the **-profile**, **addrdiscovery**, or **-lbs** parameters.
2. Use the **lshosttype** command to obtain a list of known host types.

-portgrp *port_grp_number*

(Optional) Specifies the identifier that associates two or more host ports with access to a common volume group. Port group zero is reserved for ports that 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 null (-) when you issue a **lshostconnect** or **showhostconnect** command.

-ioport *port_ID . . .*

(Optional) Specifies all, none, one, or more I/O port IDs that allow host connection access to volumes.

Enter **all** to specify that you want all I/O ports added.

Enter **none** to specify that you do not want to add I/O ports. If you do not specify I/O ports, the image 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.

I/O ports cannot share the same WWPN. Ensure that there are no conflicts with the I/O ports of existing SCSI host connections.

You can select up to 128 ports for an open systems host attachment assignment. If you enter a list of I/O port IDs, access from this host connection to the specified volume group is allowed using only the specified list.

A port ID is four hexadecimal characters in the format *EEAP*, where:

- *EE* is an I/O port enclosure number in the range of 00 - 17 (2107 machine types).
- *A* is the adapter number and is specified as 1, 2, 4, or 5 (2107 machine types).
- "*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.

-speed *host_port_speed*

(Optional) Specifies the host I/O port interface speed. This value defaults to *unknown* if it is not specified.

Note: If you want to explicitly specify a value of unknown, use 0 (zero) for the value.

You might want to specify the advertised interface speed which is generally one of the following values: 1Gb/s, 2Gb/s, 4Gb/s or 200Mb/s.

You can issue the **lshostconnect** command to view the value recognized within the system. You might find that the value specified on the report associated with the **lshostconnect** command is different than the value you have specified with this parameter. The system value overrides your input if there is a difference.

host_name | -

(Required) Specifies your host system or port name, limited to 16 characters.

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the mkhostconnect command

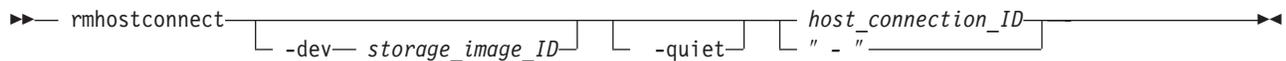
```
dscli>mkhostconnect -dev IBM.2107-75FA120 -wwname 12341234000A000F
-profile "IBM pSeries - AIX" host_1_port_1
```

The resulting output

```
Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
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 the manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for all host connections.

-quiet

(Optional) Turns off the 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 *manufacture.type-serial number/hostconnectID*) or a shortened version if the **-dev** parameter is specified.

Example of a fully qualified host connection ID: IBM.2107-75FA120/1

If you specify the dash (-), this parameter information is automatically supplied.

Attention: Use caution when you work with connection IDs to ensure that you have specified the connection that you want to delete. For instance, 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 (2107)

Invoking the **rmhostconnect** command

```
dscli>rmhostconnect -dev IBM.2107-75FA120 1
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
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—┌──────────────────────────────────┐ ┌──────────────────┐ ───────────────────►  
                    └── -dev— storage_image_ID ─┘ └── " - " ───────────────────┘
```

Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which consists of the manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the host connection.

host_connection_ID | -

(Required) Specifies a fully qualified host connection ID, which includes the manufacturer, type, and sequence 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 specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both the 1750 and 2107.

Invoking the showhostconnect command

```
dscli>showhostconnect -dev IBM.2107-75FA120
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

Name	ID	WWPN	HostType	LBS	Addrdiscovery
My host port 1	IBM.2107-75FA120/1	3007ACF30A2399E0	Unknown	512	reportLUN

Profile	Portgrp	VolgrpID	Achtopo	ESSIOport	Speed	Desc
IBM pSeries - AIX	0	100	SCSI -FCP	I0111,I0121 I0211,I0221	1 GB/s (null (-) is reported for a 1750 model)	SCSI1

Report field definitions

Name

Specifies the host connection SCSI port nickname.

The name is limited to 32 byte or 16 double-byte characters.

ID Specifies 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

Specifies the worldwide port name (WWPN) for the designated host system port.

HostType

Specifies 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

Specifies 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 iSeries host system attachments.

Addrdiscovery

Specifies 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

Specifies the name of the host connection behavior profile.

Portgrp

Specifies 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

Specifies 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.

Achtopo

Specifies the topology of the attached unit. This value is used to select which storage facility I/O ports are compatible with the attachment to the specified storage unit. One of the following values is displayed:

- Unknown
- SCSI-FCP
- FC-AL

ESSIOport

Specifies the array of port IDs that the designated host port is logged into.

The port ID component is four hexadecimal characters in the format *EEAP*, where *EE* is a port enclosure number 00 - 17 for the 2107, *A* is the adapter number 1,2,4,5 for the 2107, and *P* is the port number (0 - 3). The number is prefixed with the letter I.

Speed

Specifies the system identified host I/O port interface speed. This field displays a value of *unknown* if the interface speed was not specified. You might find that the value specified on the report is different than the value you specified with the **chhostconnect** command. The system value overrides your input if there is a difference.

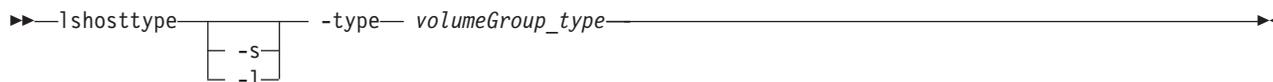
Note: This value is not reported for a 1750 model type. A null (-) value is displayed.

Desc

Specifies the description you defined for the SCSI host port. The description is limited to 256 byte or 128 double-byte characters.

Ishosttype

The **Ishosttype** command displays a list of known hosts, their associated port profiles, address discovery, and logical block size values. Use this command to get the available host types for the **mkhostconnect** command.



Parameters

- s (Optional) Displays the host types only. You cannot use the -l and -s parameters together.
- l (Optional) Displays the default output for the specified host type. You cannot use the -l and -s parameters together.
- type *volumeGroup_type*
(Required) Displays only those host types that are associated with the specified volume group type.

volumeGroup_type

Only one type can be queried at a time. The following list provides the choices that can be specified.

- ficonall
- scsiall
- scsimask
- scsimap256
- os400all
- os400mask

Example

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

The following table represents the headers that are displayed on the output report that is associated with the **lshosttype** command. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the lshosttype command

```
dscli>lshosttype -l -type
scsiall
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0

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.

Network ports

This section contains commands that are used to configure and modify network ports.

Use the following commands to configure or modify a network port for TCP/IP and to show detailed network port information. The network port commands consist of the following:

- **setnetworkport**
- **lsnetworkport**
- **shownetworkport**

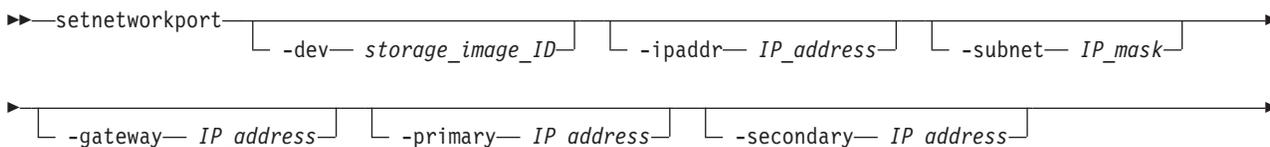
The **setnetworkport** command configures one network port for TCP/IP connections

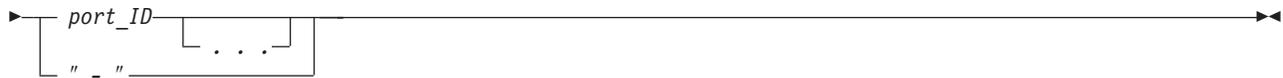
The **lsnetworkport** command generates a report that lists all the network ports that are installed in the specified storage unit.

The **shownetworkport** command generates a detailed properties report for the specified network port ID.

setnetworkport

The **setnetworkport** command (DS8000 only) configures one network port for TCP/IP connections.





Parameters

Notes:

1. One or more of the optional parameters, **-ipaddr**, **-subnet**, **-gateway**, **-primary**, or **-secondary**, must be specified.
2. Only one port can be set or modified at a time.

-dev *storage_image_ID*

(Optional) Accepts a fully qualified storage image ID. The storage image ID consists of manufacturer, machine type, and serial number.

Note: This parameter is required if you do not specify a fully qualified port ID and you do not specify a value for the *devId* variable in your profile file.

-ipaddr *IP_address*

(Optional) Specifies the TCP/IP address for the designated port.

Note: Ports that are on the same server must be on different subnets.

-subnet *IP_mask*

(Optional) Specifies the TCP/IP subnet mask for the designated port.

Note: Ports that are on the same server must be on different subnets.

-gateway *IP_address*

(Optional) Specifies the gateway TCP/IP address for the designated port.

-primary *IP_address*

(Optional) Specifies the primary DNS TCP/IP address for the designated port.

-secondary *IP_address*

(Optional) Specifies the secondary DNS TCP/IP address for the designated port.

port_ID . . . | -

(Required) Displays network ports matching the specified IDs. This parameter accepts a fully qualified port ID, which includes the storage image ID, or a shortened version without the storage image ID when the **-dev** parameter is specified.

A port ID is prefixed with letter *I* and consists of four hexadecimal characters in the format *trre eeee aaaa pppp* (value is separated for readability), where:

- *t* = port type (1= network port)
- *rr* = reserved bits and set to zero
- *eeee* = enclosure number (value specified is 18 - 1F)
- *aaaa* = adapter number
- *pppp* = port number

If you specify the dash (-) this parameter information is automatically supplied.

Example

Invoking the setioport command

This example shows how you would configure network port (IBM.2107-1300321/ I9801) to use IP address 192.555.5.5.

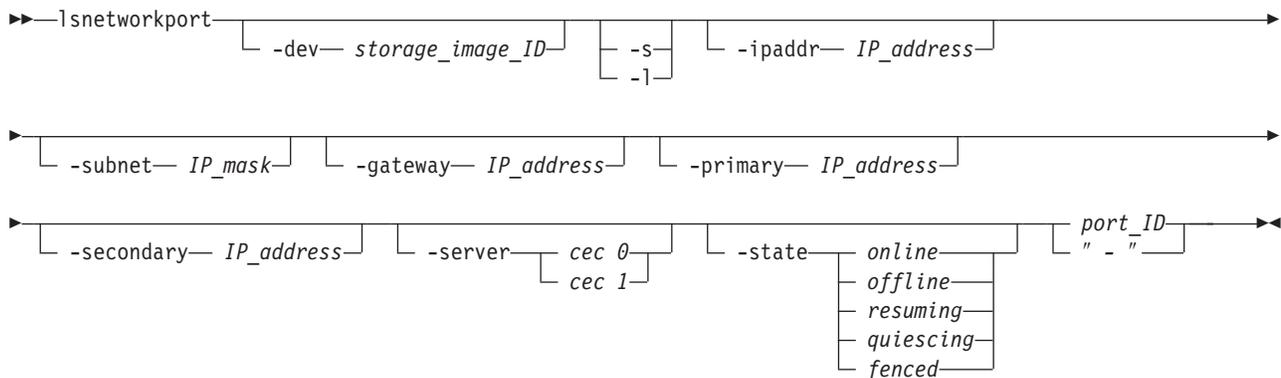
```
dscli>setnetworkport  
-dev IBM.2107-1300321 -ipaddr 192.555.5.5 I9801
```

The resulting output

```
Date/Time: January 11, 2006 236:44 PM MST DSCLI Version: 5.0.0.0  
DS: IBM.2107-1300321  
CMUC00250I setnetworkport: Network port I9801 successfully configured.
```

lsnetworkport

The **lsnetworkport** command (DS8000 only) displays a report that lists all network ports that are installed in the specified storage unit.



Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID. The storage image ID consists of manufacturer, machine type, and serial number.

Note: This parameter is required if you do not specify a fully qualified port ID and you do not specify a value for *devid* in your profile file.

-s (Optional) Specifies that only the port IDs be displayed for the designated storage unit. You cannot use the **-I** and the **-s** parameters together.

-I (Optional) Displays the default output and the values for the owning server, port speed, and port location. You cannot use the **-I** and the **-s** parameters together.

-ipaddr *IP_address*

(Optional) Specifies the TCP/IP address for the designated port.

Note: Ports that are on the same server must be on different subnets.

-subnet *IP_mask*

(Optional) Specifies the TCP/IP subnet mask for the designated port.

Note: Ports that are on the same server must be on different subnets.

- gateway** *IP_address*
(Optional) Specifies the gateway TCP/IP address for the designated port.
- primary** *IP_address*
(Optional) Specifies the primary DNS TCP/IP address for the designated port.
- secondary** *IP_address*
(Optional) Specifies the secondary DNS TCP/IP address for the designated port.
- server** *cec 0 | cec 1*
(Optional) Displays the server that is associated with the I/O port. One of the following values is displayed:
 - cec 0** Specifies that the user port is connected to processor complex 0 and that the I/O port is associated with server 0.
 - cec 1** Specifies that the user port is connected to processor complex 1 and that the I/O port is associated with server 1.
- state** *online | offline | resuming | quiescing | fenced*
(Optional) Displays network ports that are in the specified state. One of the following values is displayed:
 - online** Specifies that the network port is enabled.
 - offline**
Specifies that the network port is disabled.
 - resuming**
Specifies that the network port is in the process of becoming enabled.
 - quiescing**
Specifies that the network port is in the process of becoming disabled.
 - fenced**
Specifies that the network port is configured but there is a problem and the port cannot be used.

port_ID | -
(Required) Displays I/O ports that match the specified IDs. This parameter accepts a fully qualified port ID, which includes the storage image ID, or a shortened version without the storage image ID when the **-dev** parameter is specified.

A port ID is prefixed with letter *I* and consists of four hexadecimal characters in the format *trre eeee aaaa pppp* (value is separated for readability), where:

- *t* = port type (1= network port)
- *rr* = reserved bits and set to zero
- *eeee* = enclosure number (value specified is 18 - 1F)
- *aaaa* = adapter number
- *pppp* = port number

If you specify the dash (-), this parameter information is automatically supplied.

Example

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

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

Invoking the lsnetworkport command

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

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

ID	IP Address	Subnet Mask	Gateway	Primary DNS
IBM.2107-75FA120 /I9801	192.555.5.5	255.255.555.5	192.555.5.5	192.555.5.55

Secondary DNS	State	Server	Speed	Type	Location
192.555.5.55	Online	00	1 Gb/sec	Ethernet-Copper	U7879.001. 03AA6D0- P0-C0-T0

Report field definitions

ID Specifies the port ID.

IP Address

Specifies the TCP/IP address for the designated port.

Subnet Mask

Specifies the TCP/IP subnet mask for the designated port.

Gateway

Specifies the gateway TCP/IP address for the designated port.

Primary DNS

Specifies the primary DNS TCP/IP address for the designated port.

Secondary DNS

Specifies the secondary DNS TCP/IP address for the designated port.

State Specifies the state of the network ports. One of the following values can be displayed in this field:

online Specifies that the network port is enabled.

offline

Specifies that the network port is disabled.

resuming

Specifies that the network port is in the process of becoming enabled.

quiescing

Specifies that the network port is in the process of becoming disabled.

fenced

Specifies that the network port is configured but there is a problem and the port cannot be used.

Server Specifies which server that the I/O port is attached to.

Speed Specifies the network port interface speed.

Type Specifies the network port interface type.

Location

Specifies the network port location using the format: *Utttt.mmm.ppsssss-Pn-Cn-Tn*

- *Utttt.mmm.ppsssss* identifies the location of the processor complex.
- *Pn* identifies the planer number.
- *Cn* identifies the card number.
- *Tn* identifies the port number.

shownetworkport

The **shownetworkport** command (DS8000 only) displays detailed properties of an individual network port ID.

```
▶▶—shownetworkport— [ -dev— storage_image_ID— ] [ port_ID— ] ▶▶
```

Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID. The storage image ID consists of manufacturer, machine type, and serial number.

Note: This parameter is required if you do not specify a fully qualified port ID and you do not specify a value for *dev* in your profile file.

port_ID | -

(Required) Displays I/O ports that match the specified IDs. This parameter accepts a fully qualified port ID, which includes the storage image ID, or a shortened version without the storage image ID when the **-dev** parameter is specified.

A port ID is prefixed with letter *I* and consists of four hexadecimal characters in the format *trre eeee aaaa pppp* (value is separated for readability), where:

- *t* = port type (1= network port)
- *rr* = reserved bits and set to zero
- *eeee* = enclosure number (value specified is 18 - 1F)
- *aaaa* = adapter number
- *pppp* = port number

If you specify the dash (-), this parameter information is automatically supplied.

Example

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

The following tables represent the headers that are displayed on the output report that is associated with the **shownetworkport** command using the **-fullid** command flag.

Invoking the shownetworkport command

```
dsccli>shownetworkport
-fullid -dev IBM.2107-1300321 I9801
```

The resulting output

Date/Time: January 11, 2006 2:26:14 PM MST IBM
 DSCLI Version: 5.0.0.0 DS: IBM.2107-1300321

ID	IP Address	Subnet Mask	Gateway	Primary DNS
IBM.2107-1300321/I9801	192.555.5.5	255.255.555.5	192.555.5.5	192.555.5.55

Secondary DNS	State	Server	Speed	Type	Location
192.555.5.55	Online	00	1 Gb/sec	Ethernet-Copper	U7879.001.03AA6D0-P0-C0-T0

Report field definitions

ID Specifies the port ID.

IP Address
 Specifies the TCP/IP address for the designated port.

Subnet Mask
 Specifies the TCP/IP subnet mask for the designated port.

Gateway
 Specifies the gateway TCP/IP address for the designated port.

Primary DNS
 Specifies the primary DNS TCP/IP address for the designated port.

Secondary DNS
 Specifies the secondary DNS TCP/IP address for the designated port.

State Specifies the state of the network ports. One of the following values can be displayed in this field:

online Specifies that the network port is enabled.

offline
 Specifies that the network port is disabled.

resuming
 Specifies that the network port is in the process of becoming enabled.

quiescing
 Specifies that the network port is in the process of becoming disabled.

fenced
 Specifies that the network port is configured but there is a problem and the port cannot be used.

Server Specifies which server that the network port is attached to.

Speed Specifies the network port interface speed.

Type Specifies the network port interface type.

Location

Specifies the network port location in the format *Utttt.mmm.ppsssss-Pn-Cn-Tn* where

- *Utttt.mmm.ppsssss* identifies the location of the processor complex.
- *Pn* identifies the planer number.
- *Cn* identifies the card number.
- *Tn* identifies the port number.

Storage configuration commands

This section contains storage configuration commands.

The following commands allow you to configure storage for zSeries and open system hosts.

Array site specific commands

This section contains commands that are used to display array site information.

Use the following commands to display array site image information:

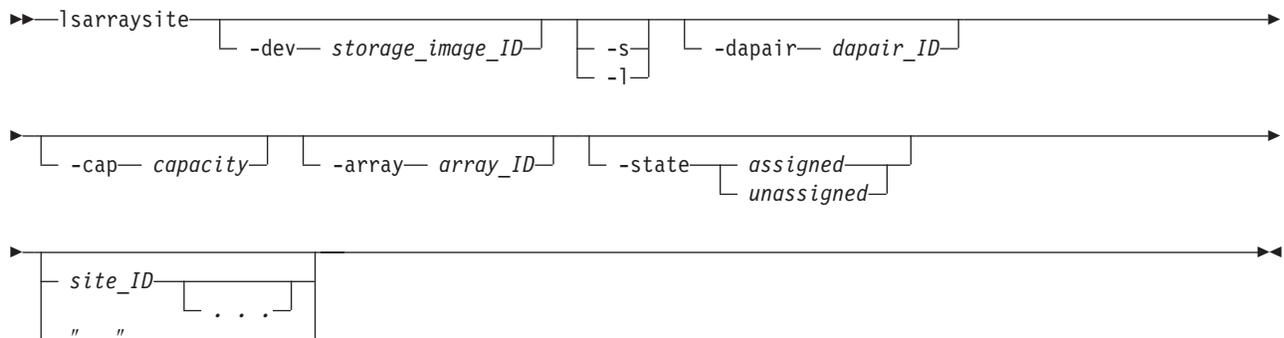
- **lsarraysite**
- **showarraysite**

The **lsarraysite** command generates a report that lists the array sites and status information for each array site in the list.

The **showarraysite** command generates a report that displays the detailed properties of a specific storage image array site.

lsarraysite

The **lsarraysite** command displays a list of array sites and status information for each array site in the list.



Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which consists of the manufacturer, type, and serial number. This parameter is required if the qualified site ID is not specified.

-s (Optional) Displays the array ID. You cannot use the **-l** and the **-s** parameters together.

-l (Optional) Displays the default output and the disk drive module rpm (revolutions per minute). You cannot use the **-l** and the **-s** parameters together.

-dapair *dapair_ID*
(Optional) Displays array sites that are associated with a common device adapter pair ID. A device adapter pair ID is a two-digit decimal number with no leading zeros.

-cap *capacity*
(Optional) Displays in 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*
(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.

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.

If you specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the **lsarraysite** command

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

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

Report field definitions

Arsite Specifies 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

Identifies 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 (10⁹ Byte)

Specifies the minimum disk capacity of the disks in the designated array site in gigabytes.

Diskrpm

Specifies the minimum disk rpm of the disks in the designated array site.

State Specifies the array site state. One of the following values can be displayed in this field:

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.

Array Specifies the array ID that this assigned array site is assigned to. The ID is prefixed by the letter *A*.

Disk class

Specifies the disk class as either high speed fibre-channel disk drives or near-line disk drives. The displayed value is one of the following:

ENT Specifies enterprise and represents high speed fibre channel disk drives.

NL Specifies near-line and represents ATA (FATA) disk drives

showarraysite

The **showarraysite** command displays detailed properties of a specific storage image array site.

```
►► showarraysite [-dev storage_image_ID] [site_ID] ◀◀
```

Parameters

-dev storage_image_ID

(Optional) Specifies the storage image ID, which consists of manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the site ID.

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 specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the showarraysite command

```
dscli>showarraysite  
-dev IBM.2107-75FA120 S11
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

arsite	DA pair	dkcap (10 ⁹ B)	Diskrpm	State
IBM.2107- 75FA120/ S11	IBM.2107- 75FA120 /0	146	15000	Assigned

Array	Dkrate (GB/sec)	DDMSN	Spares	dataDDM	Diskclass
IBM.2107- 75FA120 /A44	2	0123456789 ABCDEF	0	8	NL

Report field definitions

Arsite Specifies the array site ID. The array site ID is a four-digit decimal number, no leading zeros, prefixed by the letter S.

Notes:

1. IBM 2107 array sites consist of eight DDMs.
2. The array site ID does not imply a physical location.

DA pair

Specifies 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 (10⁹B)

Specifies the minimum disk capacity of the disks in the designated array site.

Diskrpm

Specifies the minimum disk rpm of the disks in the designated array site.

State Specifies the array site state. The values that can be displayed in this field are as follows:

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.

Array Specifies the array ID that the designated array site is assigned to. The ID is prefixed by the letter *A*.

Dkrate

Specifies the minimum disk interface rate of the disks in the designated array site.

DDMSN

Specifies 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

Specifies the number of spare DDMs that are allocated from the array site.

DataDDM

Specifies the number of data DDMs. This value is based on the number of DDMs minus the number of spares.

Disk Class

Specifies the disk class as either high speed fibre channel disk drives or near-line disk drives. The displayed value is one of the following:

ENT Specifies enterprise and designates a high-speed fibre channel disk drive.

NL Specifies near-line and represents ATA (FATA) disk drives.

Array specific commands

This section contains commands that are used to create and delete arrays and to display array information.

Use the following commands to create and delete arrays and to display array information:

- **lsarray**
- **mkarray**
- **rmarray**
- **showarray**

The **lsarray** command generates a report that displays a list of arrays in a storage image and the status information for each array in the list.

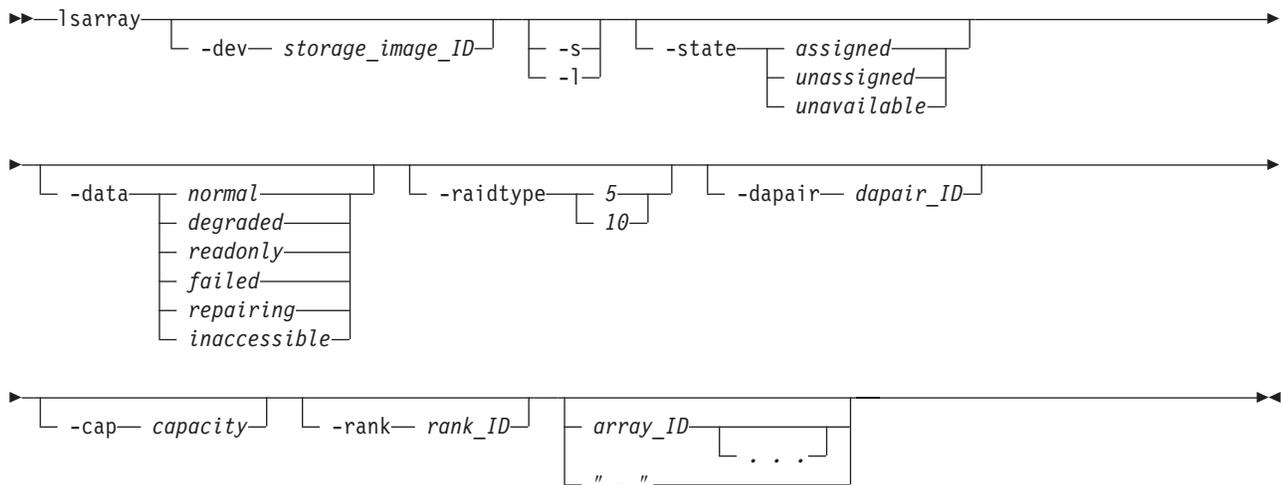
The **mkarray** command creates one array per command.

The **rmarray** command removes the specified array or arrays from the storage unit.

The **showarray** command generates a report that displays the detailed properties of a specific array.

lsarray

The **lsarray** command displays a list of arrays in a storage image and status information for each array in the list.



Parameters

- dev** *storage_image_ID*
(Optional) Specifies the storage image ID, which includes the manufacturer, type, and serial number. This parameter is required if the fully qualified array ID is not specified.
- s** (Optional) Displays the array ID. You cannot use the **-l** and the **-s** parameters together.
- l** (Optional) Displays the default output data. You cannot use the **-l** and the **-s** parameters together.
- state** *assigned* | *unassigned* | *unavailable*
(Optional) Displays only those arrays that are in the specified state.
- data** *normal* | *degraded* | *readonly* | *failed* | *repairing* | *inaccessible*
(Optional) Displays only those arrays that are in the specified data state.
- raidtype** *5* | *10*
(Optional) Displays only those arrays with the specified RAID type, 5 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 gigabytes (GB) only the array with the specified DDM capacity. You can specify up to three digits after the decimal point, for example **-cap 144.7**.
- 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*.
- 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:

```
dscli>lsarray IBM.2107-75FA120 -l A11 A12 A14-A16
```

If you specify the dash (-), this parameter information is automatically supplied..

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the lsarray command

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

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

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
S20	R11	10	145	ENT
S21	R12	11	145	ENT
S30	R13	20	300	NL
S31	-	21	300	NL

Report field descriptions

Array Specifies the array number. The array number starts with the prefix *A*, followed by a four-digit decimal number with no leading zeros (for example, A44).

State Specifies 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 or 10) and the array configuration (for example, 6+P).

arsite Indicates the array sites that are associated with the array.

Rank Specifies 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 four-digit decimal number, 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 (10⁹ Byte)

Indicates the minimum disk capacity (10⁹ Byte) of the storage devices (DDMs) in the specified array.

Disk class

Specifies the disk class as either high speed fibre-channel disk drives or near-line disk drives. The displayed value is one of the following:

ENT Specifies enterprise and represents high speed fibre-channel disk drives

NL Specifies near-line and represents ATA (FATA) disk drives

mkarray

The **mkarray** command creates one array per command.

```
mkarray -dev storage_image_ID -raidtype 5 | 10 -arsite array_site
```

Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified array site.

-raidtype *5 | 10*

(Required) Specifies a RAID type for the array.

-arsite *array_site*

(Required: 2107 explanation) Specifies the array site for the array. An array site number is a four-digit decimal number with no leading zeroes, prefixed with the letter S.

Example of fully qualified array site: IBM.2107-75FA120/S11

Example (2107)

Invoking the mkarray command

```
dscli>mkarray -dev IBM.2107-75FA120 -raidtype 10 -arsite S10
```

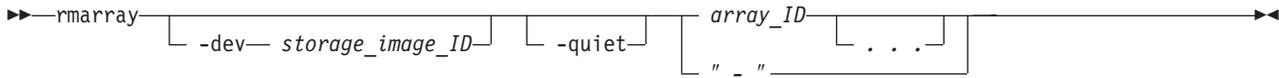
The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
Array A10 successfully created.
```

rmarray

The **rmarray** command deletes arrays.



Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for all array IDs.

-quiet

(Optional) Turns off the 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.

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the **rmarray** command

```
dscli>rmarray -dev IBM.2107-75FA120 A44
```

The resulting output

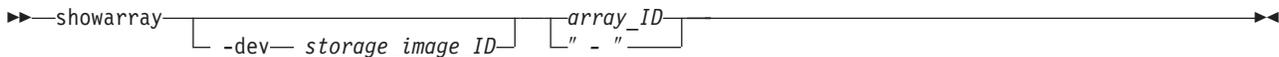
```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-68FA120
```

```
Are you sure you want to delete array IBM.2107-68FA120/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 the manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the array.

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 specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the showarray command

```
dscli>showarray -dev IBM.2107-75FA120 A44
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

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
2107-75FA123/R26	2107-75FA123/11	145	15000	FCAL	2	ENT

Report field definitions

Array Specifies the array ID number. The array ID number starts with the prefix *A*, followed by a four-digit decimal number with, no leading zeros (for example, A44).

SN Specifies the unique internal identifier for the data space of the designated array (for example, AZ123AQ).

State Specifies 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

Specifies 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 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 state 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

Specifies the type of RAID array (5 or 10) and the array configuration (for example, 6+P).

Arsite Specifies the array sites that are associated with the array.

Rank Specifies 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 four-digit decimal number, with no leading zeros (for example, R26).

Note: If the array is unassigned, the field is null (–)

DA pair

Specifies the DA pair ID. DA pairs are located 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)

Specifies the minimum disk capacity (10^9B) of the storage devices (DDMs) in the designated array.

DDMRPM

Specifies the minimum disk rpm of the DDMs in the designated array.

Interface Type

Specifies the disk interface type of the DDMs in the designated array

Interrate

Specifies the minimum disk interface rate of the disks in the designated array.

Disk class

Specifies the disk class as either high speed fibre-channel disk drives or near-line disk drives. The displayed value is one of the following:

ENT Specifies enterprise and represents high speed fibre-channel disk drives

NL Specifies near-line and represents ATA (FATA) disk drives

Rank specific commands

This section contains commands that are used to create, modify, and delete ranks and to display rank information.

Use the following commands to create, modify, and delete ranks and to display rank information:

- **chrank**
- **lsrank**
- **mkrank**
- **rmrank**
- **showrank**

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.

The **lsrank** command generates a report that displays a list of defined ranks in a storage unit and the status information for each rank in the list.

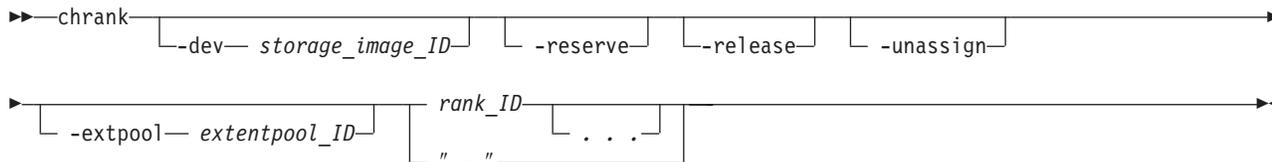
The **mkrank** command creates one fixed block or count key data (CKD) rank from one array.

The **rmrank** command deletes the specified ranks from a storage unit.

The **showrank** command generates two types of reports. One report displays the detailed properties of a specified rank. The other report displays the performance metrics of a specified rank.

chrank

The **chrank** command assigns an unassigned rank to an extent pool, or removes an assigned rank from an extent pool. This command can also be used to change an assigned rank to an unassigned rank.



Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which consists of manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for all ranks; otherwise, it is optional.

Example: IBM.2107-75FA120

-reserve

(Optional) Specifies that the rank configuration state be set to Reserved.

-release

(Optional) Specifies that the rank configuration state be set to Normal.

-unassign

(Optional) Specifies that the rank configuration state be set to Unassigned.

-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.

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.

Example: IBM.2107-75FA120/R2

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the chrank command

```
dscli>chrank -dev IBM.2107-75FA120 -extpool P101 R2
```

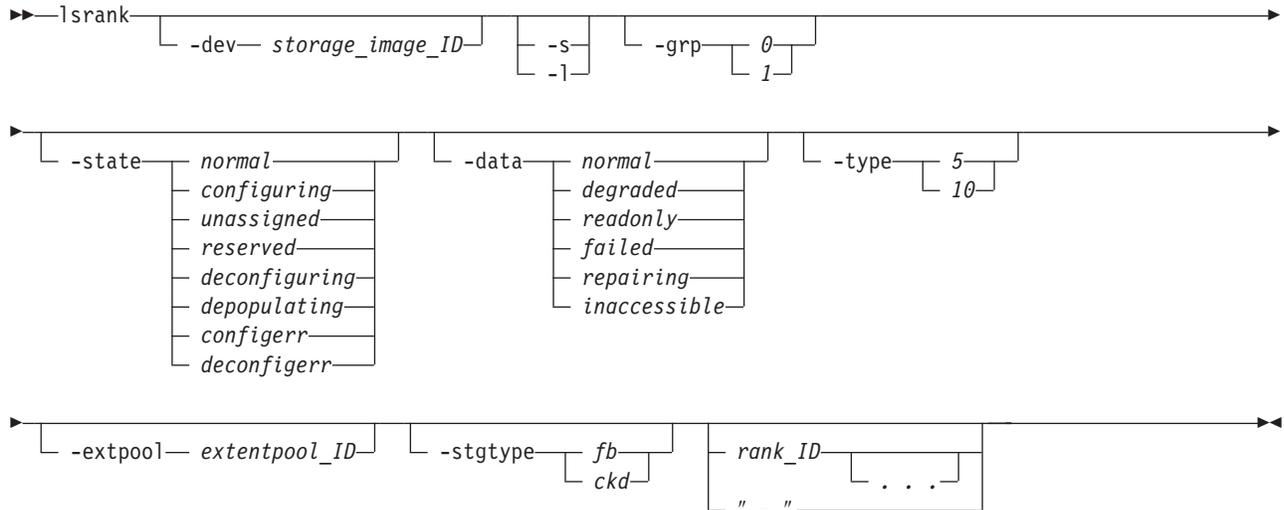
The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

Rank IBM.2107-75FA120/R2 successfully modified.

lsrank

The **lsrank** command displays a list of defined ranks in a storage image and status information for each rank.



Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, type, and serial number.

Example: IBM.2107-75FA120

-s (Optional) Displays the rank ID. You cannot use the **-l** and the **-s** parameters together.

-l (Optional) Displays the default output, extent pool name, number of extents, and extents that are used for each rank. You cannot use the **-l** and the **-s** parameters together.

-grp *0 | 1*

(Optional) Displays only the ranks that belong to the specified rank group. A rank in the unassigned state is not associated with a rank group.

-state *normal | configuring | unassigned | reserved | deconfiguring | depopulating | configerr | deconfigerr*

(Optional) Displays only ranks in the specified state.

-data *normal | degraded | readonly | failed | repairing | inaccessible*

(Optional) Displays only ranks in the specified data state.

-type *5 | 10*

(Optional) Displays only ranks of the specified RAID type.

-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 zeroes, prefixed with the letter P.

-stgtype *fb | ckd*

(Optional) Displays only ranks that are configured for the specified storage type.

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.

If you specify the dash (-), this parameter information is automatically supplied.

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 using the **-l** parameter. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the lsrank command

```
dsccli>lsrank -dev IBM.2107-75FA120 -l
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

ID	Group	State	Datstate	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

ExtentpoolID	Extpoolnam	Stgtype	Exts	Usedexts
P1	host_1_volumes	fb	1,000	500
P1	host_1_volumes	fb	1,000	500
P1	host_1_volumes	fb	1,000	500
P1	host_1_volumes	fb	1,000	500

Report field definitions

ID Specifies the unique identifier assigned to the rank.

Group Specifies the rank group that the rank is assigned to. A value of 0, 1, or - (null) is displayed. If a rank is unassigned the value displayed is - (null).

State Specifies the current configuration state of this rank ID. One of the following values is displayed:

Normal

Specifies that rank is assigned to an extent pool ID and none of the other state conditions apply.

Configuring

Specifies that the rank is in the process of being initially configured.

Unassigned

Specifies that the rank is not assigned to a extent pool ID.

Reserved

Specifies that the rank extents are not eligible for volume assignment.

Deconfiguring

Specifies that the rank is in the process of being deleted.

Depopulating

Specifies 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.

Configuration Error

Specifies 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.

Deconfiguration Error

Specifies that a rank removal process has 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.

Datstate

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).

Datstate specifies the current state of the data extents contained by the designated rank ID. One of the following values is displayed:

Normal

A rank is in the normal data state when the configuration state is one of the following: 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 Specifies the array ID that is assigned to the designated rank.

RAIDtype

Specifies the RAID type of the array associated with this rank. The value displayed is either 5 or 10.

ExtentpoolID

Specifies the extent pool to which the rank is assigned.

Extpoolnam

Specifies the name that is assigned to the extent pool to which the rank is assigned.

Stgtype

Specifies the storage type of the extent pool to which the rank is assigned. The value displayed is either fb (fixed block) or ckd (count key data)

Exts Specifies the number of extents that are contained in the designated rank. The value displayed is a number in the range of 1 - 4000.

Usedexts

Specifies the number of extents that are allocated to volumes from the designated rank. The value displayed is a number in the range of 1 - 4000.

mkrank

The **mkrank** command creates one fixed block or count key data (CKD) rank from one array.

```
mkrank [-dev storage_image_ID] -array array_ID -stgtype [fb | ckd] [-wait]
[-extpool extentpool_ID]
```

Parameters

Note: It is recommended 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 has failed for any reason.

-dev storage_image_ID

(Optional) Specifies the storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the extent pool and you do not specify a value for the *dev* variable in your profile file.

-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.

-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 (2107)

Invoking the mkrank command

```
dscli>mkrank -dev IBM.2107-75FA120 -array A44 -stgtype fb -wait
```

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 piece 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 array is listed as assigned to the rank 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 array that is assigned to the rank from which you removed it, the remove rank process 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. This parameter is required if you do not specify a fully qualified ID for all ranks; otherwise, it is optional.

-quiet

(Optional) Turns off the 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.

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the **rmrank** command

```
dscli>rmrank -dev IBM.2107-75FA120 R23
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
Are you sure you want to delete rank R23? [y/n]: Y
```

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
Rank R23 successfully deleted.
```

showrank

The **showrank** command displays detailed properties or performance metrics of a rank.

```
▶▶ showrank [-dev storage_image_ID] [-metrics] rank_ID ▶▶
```

Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which consists of manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the rank; otherwise, it is optional.

-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

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the showrank command to show rank properties

```
dscli>showrank -dev IBM.2107-75FA120 R34
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

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	Strpsize	Extsize
1	0	128	4	4	16,384

Report field definitions (-metrics parameter not specified)

ID Specifies the unique ID that is assigned by the system to the rank. The ID includes the storage image ID and the rank ID.

- SN** Specifies the unique serial number that is assigned to the designated rank ID.
- Group** Specifies the rank group that the rank is assigned to. One of the following values are displayed: 0, 1, - (null).
- Note:** Null (-) is displayed if the rank has not been assigned to an extent pool.
- State** Specifies 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**
Specifies that a rank is assigned to an extent pool ID and none of the other state conditions apply.
 - Configuring**
Specifies that a rank is in the process of being initially configured. Generally this state indicates that the associated rank transaction has not completed.
 - Unassigned**
Specifies that a rank is not assigned to an extent pool ID.
 - Reserved**
Specifies that rank extents are not eligible for volume allocation.
 - Deconfiguring**
Specifies that a rank is in the process of being deleted.
 - Configuration Error**
Specifies 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.
 - Deconfiguration Error**
Specifies 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.
- Datstate**
Specifies the current state of the data extents that are contained by this rank ID. The following values can be displayed for the rank:
- Normal**
Specifies that none of the other data states apply.
 - Degraded**
Specifies that one or more arrays in the rank are in the degraded state.
 - Read Only**
Specifies that one or more arrays in the rank are in the Read Only state.
 - Failed** Specifies that one or more arrays in the rank are in the Failed state.
 - Repairing**
Specifies that one or more arrays in the rank are in the Repairing state.

Inaccessible

Specifies that one or more arrays in the rank are in the Inaccessible state.

Array Specifies the array ID that is assigned to the designated rank.

RAIDtype

Specifies the RAID type of the array that is associated with the designated rank.

ExtpoolID

Specifies the extent pool to which the designated rank is assigned.

Extpoolnam

Specifies the extent pool to which the designated rank is assigned.

Volumes

Specifies the volume IDs that have an extent pool value that is allocated on the designated rank.

Stgtype

Specifies the storage type of the extent pool the designated rank is assigned to. Valid values are fixed block and count key data (CKD).

Exts Specifies the number of extents that are contained in the designated rank. 1 - 4000 are valid values.

Usedexts

Specifies the number of extents that are allocated to volumes from the designated rank.

Widearrays

Specifies the number of wide arrays that are contained by the designated rank. 0 or 1 are valid values.

Nararrays

Specifies the number of narrow arrays that are contained by the designated rank.

Trksize

Specifies 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

Specifies the number of logical tracks in a strip on the designated rank.

Strpesize

Specifies the number of logical tracks in a stripe on the designated rank.

Extsize

Specifies the number of logical tracks in an extent on the designated rank.

Notes:

1. A CKD 1 GB extent contains 16 696 tracks.
2. A fixed block 1 GB extent contains 16 384 tracks.

Performance request

Invoking the showrank command to show performance metrics

```
dscli>showrank -dev IBM.2107-75FA120 - metrics R34
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

ID	Date	Byteread	Bytewrite	Reads
2107 -75FA120 /R34	10/11/04 02:23:47	10000	10000	10000

Writes	Timeread	Timewrite
10000	10000	10000

Report field definitions (with the -metrics parameter specified)

ID Specifies the unique ID that is assigned by the system to the rank. The ID includes the storage image ID and the rank ID.

Date Specifies the time stamp for the rank performance counters.

Byteread

Specifies the number of rank bytes that are read in 128 KB increments.

Bytewrite

Specifies the number of rank bytes that are written in 128 KB increments.

Reads Specifies the rank read operations.

Writes Specifies the rank write operations.

Timeread

Specifies the rank read response time in 16 millisecond increments.

Timewrite

Specifies the rank read response time in 16 millisecond increments.

Extent pool specific commands

This section contains commands that are used to create, modify, and delete extent pools and to display extent pool information.

Use the following commands to create, modify, and delete extent pools and to display extent pool information:

- **chextpool**
- **lsextpool**
- **mkextpool**
- **rmextpool**
- **showextpool**

The **chextpool** command modifies an extent pool.

The **lsextpool** command 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.

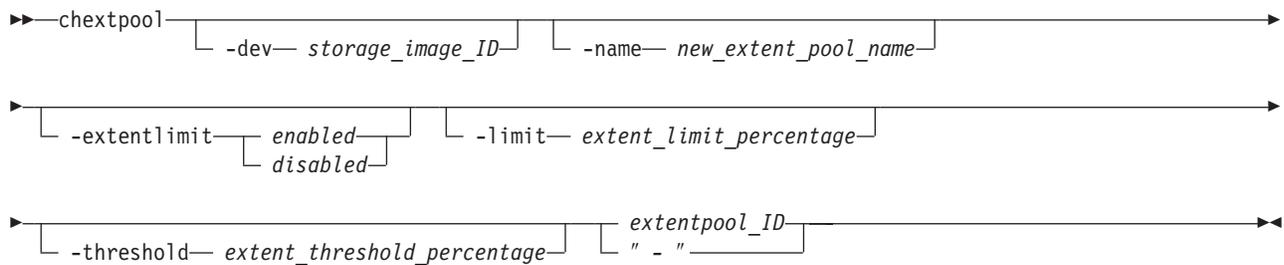
The **mkextpool** command creates a fixed block or count key data (CKD) storage type extent pool.

The **rmextpool** command deletes one or more specified extent pools from a storage unit.

The **showextpool** command 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

The **chextpool** command modifies an extent pool.



Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which consists of manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the extent pool; otherwise, it is optional.

-name *new_extent_pool_name*

(Optional) Specifies a new name for the extent pool.

Note: To change any other extent pool attributes, you must delete the extent pool and create a new one.

-extentlimit *enabled | disabled*

(Optional) Specifies that the extent limit function 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.

extentpool_ID | -

(Required) Specifies the ID of the extent pool to be changed. 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 specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the chextpool command

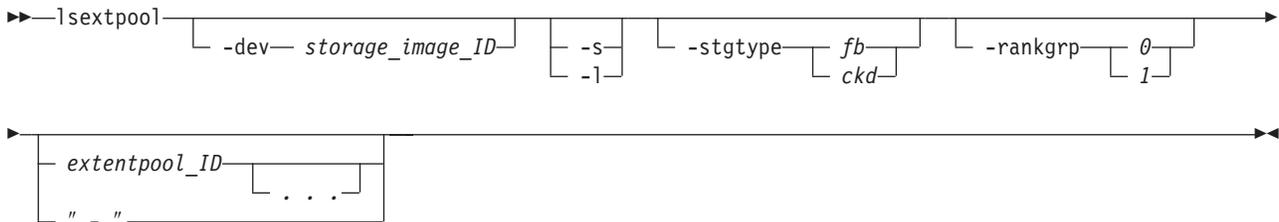
```
dsccli>chextpool -name host_4_extpool IBM.2107-75FA120/P21
```

The resulting output

Extent Pool IBM.2107-75FA120/P21 successfully modified.

lsextpool

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, type, and serial number.
- s** (Optional) Displays extent pool IDs. You cannot use the **-l** and the **-s** parameters together.
- l** (Optional) Displays default output plus additional attributes that are identified as long output. You cannot use the **-l** and the **-s** parameters together.
- stgtype** *fb | ckd*
(Optional) Displays only extent pools with the specified storage type.
- rankgrp** *0 | 1*
(Optional) Displays only extent pools in the specified rank group.
- 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.

If you specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 as the information is the same for both the 2107 and 1750.

Invoking the lsextpool command

```
dsccli>lsextpool -dev IBM.2107-75FA120 -1
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

Name	ID	Stgtype	Rankgrp	Status	Avail-stor (2^30B)
host_1_extpool	IBM.2107-75FA120/P21	fb	0	exceeded	1000
host_1_extpool	IBM.2107-75FA120/P22	fb	1	exceeded	1000
host_2_extpool	IBM.2107-75FA120/P23	fb	0	exceeded	1000
host_2_extpool	IBM.2107-75FA120/P24	fb	1	exceeded	1000
host_3_extpool	IBM.2107-75FA120/P25	fb	0	exceeded	1000
host_3_extpool	IBM.2107-75FA120/P26	fb	1	exceeded	1000

%allo-cated	Avail-able	Reser-ved	Num-vols	Num-ranks
10	1000	100	4	4
10	1000	100	4	4
10	1000	100	4	4
10	1000	100	4	4
10	1000	100	4	4
10	1000	100	4	4

Report field definitions

Name Identifies the name you assigned to the extent pool.

ID Specifies the system assigned unique identifier for the extent pool object.

stgtype

Identifies the storage type associated with the extent pool. One of the following is displayed:

- fb
- ckd

rankgrp

Specifies the rank group in which the designated extent pool is configured.

Status Specifies the extent status. One of the following values is displayed:

exceeded

Specifies that the %Extents available is greater than the extent threshold

below

Specifies that the %Extents Available is less than the extent threshold

full

Specifies that the %Extents Available is 0.

availstor (2^30 Bytes)

Specifies the available storage for the designated extent pool, in gigabytes (2 ^ 30 Bytes).

%allocated

Specifies the percentage of extents allocated.

available

Specifies the maximum number of extents available for allocation in the designated extent pool.

reserved

Specifies the extents reserved in the designated extent pool.

numvols

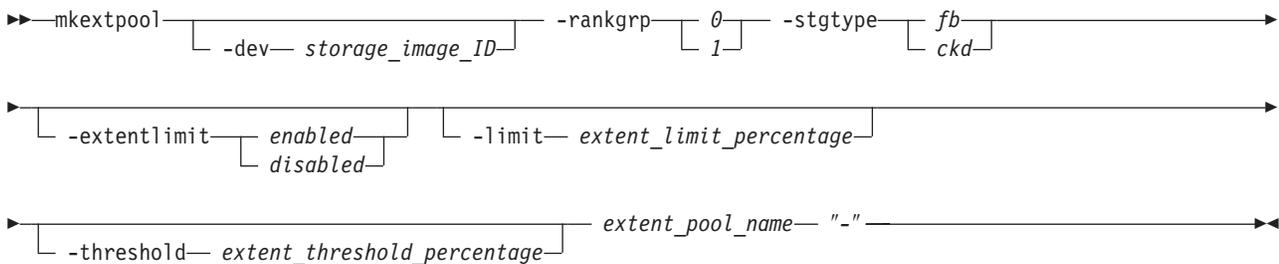
Identifies the number of logical volumes that have been configured from the designated extent pool.

numranks

Identifies the number of ranks that have been configured in the designated extent pool.

mkextpool

The **mkextpool** command creates a fixed block or count key data (CKD) storage type extent pool.

**Parameters****Notes:**

1. 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.
2. Create extent pool objects before creating array and rank objects.
3. 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.

- A server has failed and the failover and failback sequence is performed.

extentpool_ID | -

(Required) Specifies the extent pool to be displayed. 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 zeroes, 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 specify the dash (-), this parameter information is automatically supplied.

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 **showextpool** command. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the showextpool command to show extent pool properties

```
dscli>showextpool -dev IBM.2107-75FA120 P101
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

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	70

Report field definitions

Name Identifies the name you assigned to the extent pool.

ID Specifies the system assigned unique identifier for the extent pool object.

stgtype

Identifies the storage type associated with the extent pool. One of the following is displayed:

- fb
- ckd

totlstor (2^30 Bytes)

Specifies the amount of storage associated with the extent pool, in gigabytes.

availstor (2^30 Bytes)

Specifies the available storage for the designated extent pool, in gigabytes.

resvdstor (2^30 Bytes)

Specifies the amount of reserved storage for the designated extent pool, in gigabytes.

rankgrp

Specifies the rank group in which the designated extent pool is configured.

numranks

Specifies the number of ranks configured in the designated extent pool.

numvols

Identifies the number of logical volumes that have been configured from the designated extent pool.

status Specifies the extent status. One of the following values is displayed:

exceeded

Specifies that the %Extents available is greater than the extent threshold

below

Specifies that the %Extents Available is less than the extent threshold

full

Specifies that the %Extents Available is zero.

%allocated

Specifies the percentage of extents allocated. A value of 1 - 100 is displayed.

%available

Specifies the percentage of extents that are available. A value of 1 - 100 is displayed.

configured

Specifies the number of extents that are contained in the extent pool.

allowed

Specifies the number of extents that are below the applicable extent limit.

available

The number of extents of a given type that are available for allocation to a logical volume.

allocated

Specifies the number of extents of a given type in the extent pool that are allocated to logical volumes or auxiliary volumes.

reserved

Specifies the number of unallocated extents of a given type in the extent pool that are on ranks of the same extent type in the reserved state. In addition, this value includes the number of unallocated 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 threshold as a percentage of the real extents that are available compared with the actual percentage of real extents that are available.

Performance request**Invoking the showextpool command to show performance metrics**

```
dscli>showextpool -metrics IBM.2107-75FA120/P101
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

ID	Date	real extcap	realext	real allocext	real extconv
IBM.2107- 75FA120 /P101	10/11/04 02:23:47	10000	10000	10000	10000

dyreloc source	dyreloc target
10000	10000

Report field definitions

ID Specifies the system assigned unique identifier for the extent pool object.

Date Identifies the current time stamp for the extent pool performance counters.

realextcap

Specifies the real extent pool capacity in gigabytes.

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

Address group specific commands

This section contains commands that are used to display address group information.

Use the **lsaddressgrp** command to display address group information.

The **lsaddressgrp** command generates a report that displays a list of address groups for a storage unit and the status information for each address group in the list.

lsaddressgrp

The **lsaddressgrp** command displays a list of address groups for a storage image and the status information for each address group in the list.

```
lsaddressgrp [-dev storage_image_ID] [-s] [-l] [-stgtype fb | ckd]
```

Parameters

- dev** *storage_image_ID*
(Optional). Specifies the storage image ID, which consists of manufacturer, type, and serial number. Displays only the objects for the storage unit specified.
Example: IBM.2107-75FA120
- s** (Optional). Displays the address group IDs only. You cannot use the **-l** and the **-s** parameters together.
- l** (Optional). Displays the default output. You cannot use the **-l** and the **-s** parameters together.
- stgtype** *fb* | *ckd*
(Optional). Displays only the address groups that are associated with the specified storage type.

Example

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 **lsaddressgrp** command using the **-l** parameter. A separate example is not shown for the 1750 as the information is the same for both the 2107 and 1750.

Invoking the lsaddressgrp command

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

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

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

ID	Stgtype	Basevolnum	Vols	LSSs	Confgvols
IBM.2107 -75FA120 /3	ckd	0300	4096	16	164096

Report field descriptions

ID Specifies the storage image address group unique identifier. The address number is a single hexadecimal character (0 - 9 or uppercase A - F).

Stgtype

Specifies the type of logical devices that are configured for the specified address group: fb - fixed block and ckd - count key data.

Basevolnum

Specifies the lowest logical volume number in the address group.

Vols Specifies the number of logical volume numbers that are associated with the address group

LSSs Specifies the number of logical subsystems (LSSs) that are configured on the address group.

Confgvols

Specifies the number of logical volumes that are configured on the address group.

Logical control unit specific commands

This section contains commands that are used to create, modify, and delete logical control units for zSeries systems and to display logical control unit information.

Use the following commands to create, modify, and delete logical control units and to display logical control unit information:

- **chlcu**
- **lslcu**
- **mklcu**
- **rmlcu**
- **showlcu**

The **chlcu** command modifies a logical control unit.

The **lslcu** command 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.

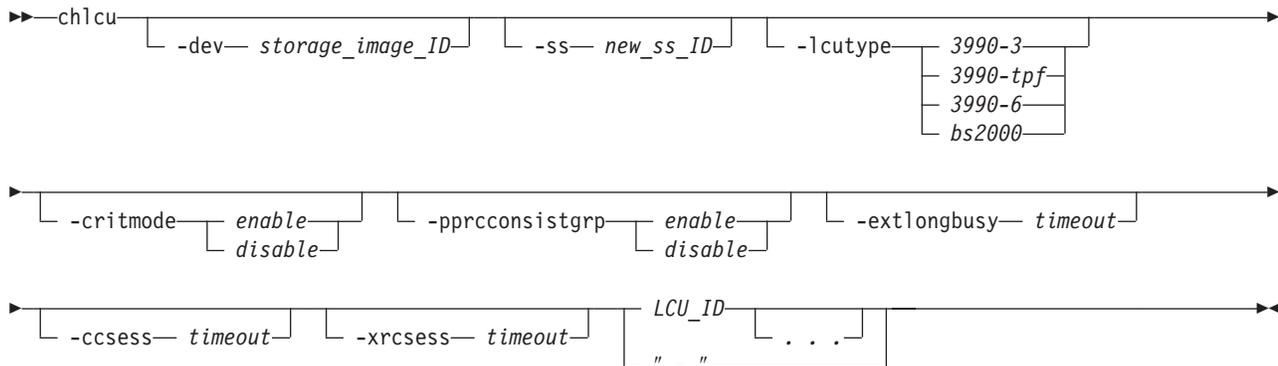
The **mklcu** command creates a logical control unit in a storage unit. A logical control unit is configured to represent a grouping of logical CKD volumes.

The **rmlcu** command deletes one or more specified logical control units.

The **showlcu** command generates a report that displays the detailed properties for the specified logical control unit.

chlcu

The **chlcu** command modifies a logical control unit.



Parameters

-dev *storage_image_ID*

(Optional). Specifies the storage image ID, which consists of manufacturer, type, and serial number.

Example: IBM.2107-75FA120

-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*

(Optional). 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.

-pprconsistgrp *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, 2107 only). 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

LCU_ID . . . | -

(Required). Specifies one or more LCUs that are to be modified by this command. An LCU ID is two hexadecimal characters 00 - FE. You must separate multiple IDs and multiple ID ranges with a space. Accepts a fully qualified LCU 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.

If you specify the dash (-), this parameter information is automatically supplied.

Example: 00-03 08

Example (2107)

Invoking the chlcu command

```
dscli>chlcu -dev IBM.2107-75FA120 -critmode enable 00-0F
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
LCU 00 successfully modified.
```

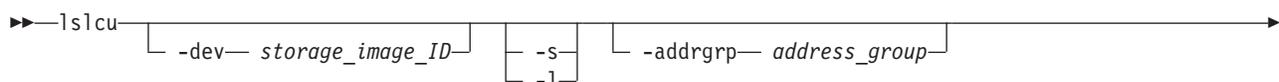
```
LCU 01 successfully modified.
```

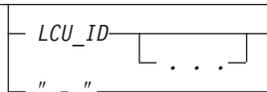
```
...
```

```
LCU 0F successfully modified.
```

lslcu

The **lslcu** command displays a list of logical control units (LCUs) for a storage image and status information for each logical control unit in the list.





Parameters

-dev *storage_image_ID*

(Optional). Specifies the storage image ID, which consists of manufacturer, type, and serial number. Displays only the objects for the storage unit that is specified.

Example: IBM.2107-75FA120

-s (Optional). Use this parameter to display LCU IDs only. You cannot use the **-l** and the **-s** parameters together.

-l (Optional). Use this parameter to display the default output. You cannot use the **-l** and the **-s** parameters together.

-addrgrp *address_group*

(Optional). Specifies an address group. Only the LCUs that belong to the specified address group are displayed. An address group is a single character in the range of 0 - 9 or A - F.

LCU_ID . . . | -

(Optional). Specifies the ID associated with an LCU. An LCU ID is two hexadecimal characters 00 - FE (2107 only).

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.

If you specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the lslcu command

```
dsccli>lslcu -dev IBM.2107-75FA120 -l
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM
DS CLI Version: 5.0.0.0 DS:
IBM.2107-75FA120
```

ID	Group	addr-grp	confg-vols	subsys	con-base-type
IBM.2107-75FA120/10	0	0	256	8010	3990-6
IBM.2107-75FA120/11	1	0	256	8011	3990-6
IBM.2107-75FA120/12	0	0	256	8012	3990-6
IBM.2107-75FA120/13	1	0	256	8013	3990-6

Report field definitions

ID Specifies 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 Specifies the server that is managing the logical control unit group. The server is identified as either 0 or 1.

Addrgrp

Specifies the address group object of which the logical control unit is a member.

Confgvols

Specifies the number of volumes, or logical devices assigned to the LCU ID. This number includes base CKD volumes and alias CKD volumes.

Subsys

Specifies the value you assigned, or the default SSID value.

Conbasetype

Specifies the LCU type. The allowable values include the following LCU types:

- 3390-3
- 3390-tpf
- 3390-6 (this is the default value if no value is assigned)
- bs2000

mklcu

The **mklcu** command creates a logical control unit (LCU) in a storage image.

```

▶▶ mklcu [-dev storage_image_ID] -qty quantity -id lcu_ID -ss ss_ID ▶▶

```


-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-tp
type 3990 model 3 for tpf

3990-6 type 3990 model 6

bs2000
type bs 2000

-critmode

(Optional). 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.

-pprconsistgrp

(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. 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. 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. Default value is 300 seconds.

Example: 500

Example (2107)

Invoking the **mklcu** command

```
dscli>mklcu -dev IBM.2107-75FA120 -qty 16 -id 80 -ss 2300
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
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, type, and serial number. This parameter is required if you do not specify a fully qualified ID for all logical control units.

Example: IBM.2107-75FA120

-quiet

(Optional). Turns off the 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 (2107 only).

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

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the **rmlcu** command

```
dscli>rmlcu -dev IBM.2107-75FA120 00-0F
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
Are you sure you want to delete LCU 00 ? y/n Y
```

```
LCU 00 successfully deleted.
```

```
Are you sure you want to delete LCU 01 ? y/n Y
```

```
LCU 01 successfully deleted.
```

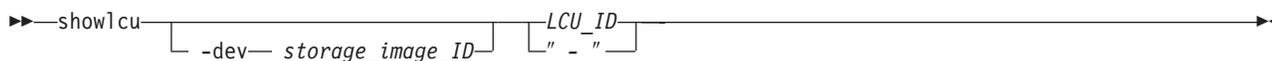
```
...
```

```
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, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the logical control unit.

Example: IBM.2107-75FA120

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 (2107 only) 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 specify the dash (-), this parameter information is automatically supplied.

Example: IBM.2107-75FA120/10

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 **showlcu** command. A separate example is not shown for the 1750. However, there is one difference between what is displayed on the output report for a 2107 versus a 1750. The 2107 displays a value for `xrcsesssttimeout` (XRC session timeout), which is not supported on the 1750.

Invoking the showlcu command

```
dscli>showlcu -dev IBM.2107-75FA120 10
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

ID	Group	Addr-grp	Config-vols	Subsys	Con-base-type
IBM.2107-75FA120/10	0	1	256	0010	3990-6

pprc-consist-grp	xtndlbz-timeout (secs)	ccsess-timeout (secs)	xrc-sess-timeout (secs)	crit-hvmode
Disabled	120	300	(2107 only) 300	Disabled

Report field definitions

ID Specifies the unique identifier that is assigned to this logical control unit object. It includes the storage image ID and the 2-digit LCU ID 2107 only (00 - FE)

Group Specifies the server that manages the logical control unit group. The displayed values are 0 or 1.

Addrgrp

Specifies the address group object that the logical control unit is a member of.

Confgvols

Specifies 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

Specifies the subsystem ID that you assigned to this logical control unit. The range of values that you selected from is 0001 - FFFF.

Conbasetype

Specifies 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.

Pprconsistgrp

Specifies the assigned PPRC consistency group state setting. If you do not designate enabled, the default value of disabled is assigned.

Xtnldbztimeout (secs)

Specifies the assigned extended long busy timeout value. If you do not designate a value, the default value of 210 seconds is assigned and displayed.

Ccsesstimout (secs)

Specifies 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 (secs)

Specifies the assigned XRC session timeout value. If you do not designate a value, the default value of 300 seconds is assigned and displayed.

Crithvmode

Specifies 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.

CKD logical volume specific commands

This section contains commands that are used to create, modify, and delete count key data (CKD) logical volumes for zSeries systems and to display logical volume information.

Use the following commands to create, modify, and delete count key data logical volumes and to display logical volume information:

- **chckdvol**
- **lsckdvol**
- **mkaliasvol**

- **mkckdvol**
- **rmckdvol**
- **showckdvol**

The **chckdvol** command modifies the nickname that you assigned to the count key data (CKD) base volume.

The **lsckdvol** command 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.

The **mkaliasvol** command creates zSeries CKD alias volumes (generally referred to as parallel access volumes or PAVs) in a storage unit.

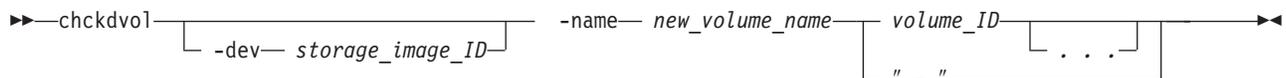
The **mkckdvol** command creates zSeries count key data (CKD) base CKD volumes in a storage image.

The **rmckdvol** command deletes one or more specified count key data (CKD) base or alias volumes from a storage unit.

The **showckdvol** command 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 changes the name of a count key data (CKD) base volume.



Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified *volume_ID*.

-name *new_volume_name*

(Required) User specified nickname for this CKD base volume. This nickname should not exceed 16 characters. It can contain one of the following wild cards:

- (#d) - insert *volume_ID* (decimal format)
- (#h) - insert *volume_ID* (hexadecimal format)

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 hexadecimal *LLVV*, where *LL* is a logical control unit number (2107 only) 00 - FE , and *VV* is a volume number (00 - FF) that is

contained by a logical subsystem object. You must fully qualify the volume ID with manufacturer, type, and serial number if you do not use the **-dev** parameter.

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the `chckdvol` command

```
dscli>chckdvol -dev IBM.2107-75FA120 -name my_volume_#d 0100
```

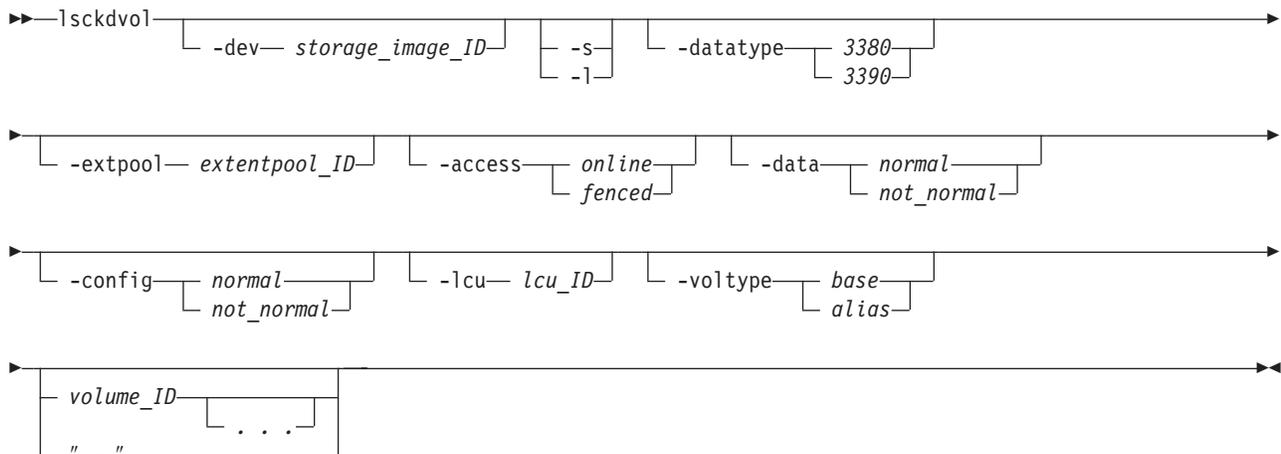
The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
CKD volume 0100 successfully modified.
```

`lsckdvol`

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

Note: For a storage unit that is heavily configured, it is recommended that you include the **-lcu** parameter as part of your command.

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, type, and serial number.

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

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

-datatype 3380 | 3390

(Optional) Displays only volumes of the specified volume data type.

-extpool *extentpool_ID*

(Optional) Display only volumes that are associated with the specified extent pool.

-access *online* | *fenced*
(Optional) Displays only volumes with the specified access state.

-data *normal* | *not_normal*
(Optional) Displays only volumes with the specified data state.

-config *normal* | *not_normal*
(Optional) Displays volumes with the specified configuration state.

-lcu *lcu_ID*
(Optional) Displays only volumes with IDs that contain the specified logical control unit ID. Each logical control unit can contain up to 256 volumes. A logical control unit ID is two hexadecimal characters 00 - FE (for 2107).

Note: If you are doing a query on a fully configured storage unit, it is recommended that you specify a specific LCU. If you do not specify a specific LCU, the entire storage unit is queried which results in a longer processing time.

-voltype *base* | *alias*
Specifies the type of CKD volumes you want displayed.

volume_ID . . . | -
(Optional) Displays volumes with the specified IDs. A volume ID is four hexadecimal characters in the range 0000 - FEFF (for 2107).

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

If you specify a dash (-), this parameter information is automatically supplied.

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 **lsckdvol** command using the **-l** parameter. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

The following example is based on the output results for a volume with 3340 cylinders.

Invoking the lsckdvol

```
dscli>lsckdvol  
-dev IBM.2107-1300861 -l 1410
```

The resulting output

```
Date/Time: April 12, 2005 6:06:31 PM JST  
IBM DS CLI Version: 5.0.3.40  
DS: IBM.2107-1300861
```

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	extentpool	sam	cap (cyl)	cap (2^30B)	cap (10^9B)
CKD base	-	P2	standard	3340	2.6	2.8

Report field definitions

Name Identifies the nickname that you assigned to the designated volume object.

ID Identifies the unique identifier that is assigned to the designated volume object

Accstate (access state)

One of the following designations can be displayed:

Online

Specifies that the logical volume is accessible to a host.

Fenced

Specifies that the logical volume is in the volume fenced state and is not accessible to the host.

- (null)

Specifies that the logical volume is designated as a CKD alias volume.

Datstate

One of the following designations can be displayed:

Normal

Specifies that none of the other data states apply. The access state is Online.

Pinned

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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.

- (null)

Specifies 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.

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.

Migrating

Indicates that the logical volume is in the process of performing dynamic volume relocation to a specified extent pool.

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.

Corrective action: Use the **rmfbvol** command to delete each volume 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.

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.

deviceMTM

One of the following is displayed:

- 3380-2
- 3380-3
- 3390-3
- 3390-9

Volume MTM is determined by the CKD base volume data type and volume capacity (in cylinders).

Volser Specifies the base CKD volume serial number written by the user at track address 0x0000, record 3.

Datatype

Identifies the volume data type setting.

Voltype

Specifies that the logical volume is one of the following:

- CKD base
- CKD alias

Orgbvols (original base vol)

One of the following will be specified:

- Identifies the original base CKD volume ID to which this CKD alias volume is assigned.
- For a CKD Alias type volume, the base and alias volume IDs are of a common LCU ID.

- - (null) is displayed for a CKD base volume type.

Extpool

Identifies the extent pool ID. Volume extents are allocated from this extent pool ID.

SAM (Standard Allocation Method)

Specifies the space efficient implementation. The current designation is Standard.

Cap (cyl)

Specifies the quantity of volume CKD cylinders that are available for host system access.

Cap (2³⁰B)

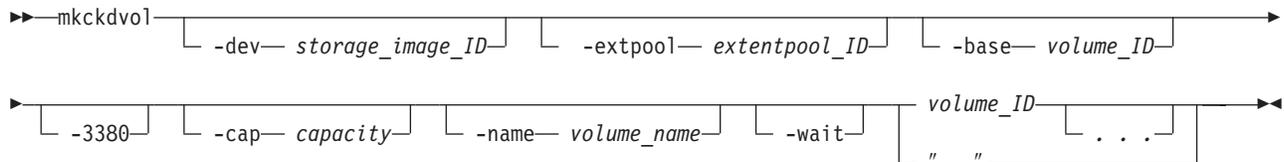
Specifies the size of volume that is available for host system access in 2³⁰B (binary GB) unit.

Cap (10⁹B)

Specifies the size of volume that is available for host system access in 10⁹B (decimal GB) unit.

mkckdvol

The **mkckdvol** command creates zSeries 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, type, and serial number.

-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 not required and 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*, where *LL* is a logical control unit number 00 - FE (2107 only), and *VV* is a 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: It is recommended that you use the **-mkaliasvol** command to create alias volumes.

-3380

(Optional) Specifies the 3380 volume type setting. Otherwise, the 3390 volume type is created. The **-3380** parameter is ignored when the **-base** parameter is specified.

Note: You must ensure that the volume data type is compatible with the host systems that access this volume.

-cap *capacity*

(Optional) Specifies the quantity of CKD cylinders that are allocated to this volume. The **-cap** parameter is ignored when the **-base** parameter is specified.

For a 3380 specification, the **-cap** parameter value is 2226 (1.59 GB) or 3339 (2.37 GB).

For 3390 volumes (default type), the **-cap** parameter can be a value in the range of 1 to 65 520 (849KB to 55.68 GB).

Note: This parameter is not required, and it is ignored if the **-base** parameter is specified.

-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.

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 hexadecimal *LLVV* where *LL* is a logical control unit number 00 - FE (2107 only) , and *VV* is a 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

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the mkckdvol command

```
dscli>mkckdvol -dev IBM.2107-75FA120
-extpool P1 -name my_volume_#d -cap 3339 0100 0101 0102 0103
```

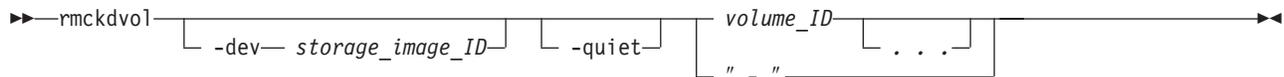
The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS IBM.2107-75FA120
```

```
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

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 includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for all logical volumes.

-quiet

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

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 hexadecimal *LLVV* where *LL* is a logical control unit (LCU) number 00 - FE (2107 only), and *VV* is a volume number (00 - FF) that is contained by a logical control unit (logical subsystem).

Note: A CKD base volume cannot be removed if an alias volume is associated with it.

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.

(2107) Example: 0100-010F 0180-FEFF 0120

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the rmckdvol command

```
dscli>rmckdvol -dev IBM.2107-75FA120 0000 0001
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
Are you sure you want to delete Volume 0000? y/n Y
```

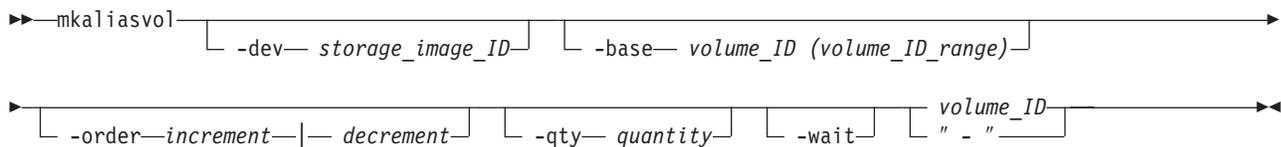
```
Volume 0000 successfully deleted.
```

```
Are you sure you want to delete Volume 0001? y/n Y
```

```
Volume 0001 successfully deleted.
```

mkaliasvol

The **mkaliasvol** command creates zSeries CKD alias volumes (generally 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 includes manufacturer, type, and serial number.

-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 will be assigned to each specified CKD base volume.

-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 hexadecimal *LLVV* where *LL* is a logical control unit number 00 - FE (2107 only) , and *VV* is a volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the mkaliasvol command

```

dscli> mkaliasvol -dev IBM.2107-75FA120 -base 0100-010F -order
decrement -qty 2 01FF

```

The resulting output

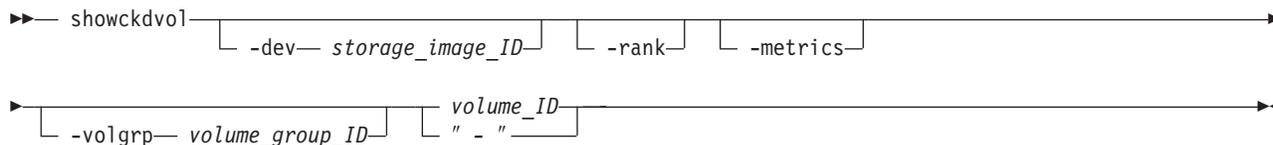
```

Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version:
5.0.0.0 DS: IBM.2107-75FA120
CKD Volume 01FF successfully created.
CKD Volume 01FE successfully created.
...
CKD Volume 01E1 successfully created.
CKD Volume 01E0 successfully created.

```

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.



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.

- There has been a server failure and the server failover or failback sequence has been initiated.

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which consists of the manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the logical volume.

-rank

(Optional) Specifies that a rank extents table 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: This parameter cannot be used with the **-metrics** or **-volgrp** parameters.

-metrics

(Optional) Displays the volume ID and performance metrics for the specified volume.

Note:

- 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.
- Do not use this parameter with the **-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 be displayed. There is only one volume group for CKD volumes and it contains all volumes.

Notes:

1. The **-volgrp** parameter can only be used 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 **-rank** or **-metrics** parameters.

volume_ID | -

(Required when 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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE, and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 **showckdvol** 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.

A separate example is not shown for the 1750 because the information is the same for both machine types. The only difference is the machine type designations 2107 versus 1750.

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.

```
dscli>showckdvol -dev
IBM.2107-1300861 -rank 1410
```

The resulting output

```
Date/Time: April 13, 2005 3:44:27 PM JST IBM DS CLI Version: 5.0.3.40 DS:
IBM.2107-1300861
```

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	3

Rank	Extents
R0	1
R2	2

Report field definitions (*-metrics* parameter not specified)

Name

Specifies the nickname that you assigned for this volume object.

ID Specifies the unique identifier that is assigned to this volume object.

Accstate

One of the following designations can be displayed:

Online

Specifies that the logical volume is accessible to a host.

Fenced

Specifies that the logical volume is in the volume fenced state and is not accessible to the host.

- (null)

Specifies that the logical volume is designated as a CKD alias volume.

Datstate

One of the following designations can be displayed:

Normal

Specifies that none of the other data states apply. The access state is Online.

Pinned

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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

Specifies 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.

– **(null)**

Specifies 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.

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.

Migrating

Indicates that the logical volume is in the process of performing dynamic volume relocation to a specified extent pool.

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.

Corrective action: Use the **rmfbvol** command to delete each volume 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.

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.

deviceMTM

One of the following will be displayed:

- 3380-2
- 3380-3
- 3390-3
- 3390-9

Volser

Specifies the volume serial number. A volume serial number is six bytes of data, displayed as six characters.

Datatype

Specifies the volume data type setting (3380 or 3390).

Voltype

Specifies that the logical volume is one of the following:

- CKD base
- CKD alias

Orgbvols(original base vol)

One of the following will be specified:

- Identifies the original base CKD volume ID to which this CKD alias volume is assigned.
- For a CKD Alias type volume, the base and alias volume IDs share a common LCU ID.
- - (null) is displayed for a CKD base volume type.

Addrgrp

Specifies the address group that contains this volume object. An address group ID is one hexadecimal character (0 - F).

Extpool

Specifies the extent pool ID.

Exts

Specifies the number of extents used by the designated volume ID.

Cap (cyl)

Specifies the quantity of volume cylinders that are available for host system access.

Cap (2³⁰B)

Specifies the size of volume that is available for host system access in (2³⁰B binary GB) units.

Cap (10⁹B)

Specifies the size of volume that is available for host system access in (10⁹B decimal GB) units.

Ranks

Specifies the number of ranks the volume resides on.

Rank (Rank Extent table)

Specifies the rank ID.

Extents (Rank Extents table)

Specifies 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 using the **-metrics** parameter.

Invoking the showckdvol to show performance metrics

```
dscli>showckdvol -metrics
IBM.2107-75FA120/0101
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI
Version: 5.0.0.0 DS: IBM.2107-75FA120
```

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	cachfwr- reqs	cachfwr- 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
10000	10000	10000

Report field definitions (*-metrics* parameter specified)

ID Specifies the unique identifier that is assigned to this volume object.

Date

Specifies the current time stamp for the volume performance counters.

normrdrqts

Specifies the number of normal read operations issued by a host to a volume.

normrdhits

Specifies the number of normal read operations where data did not move to or from a storage device.

normwritereq
Specifies Write Normal I/O Requests

normwritehits
Specifies DASD Fast Write I/O Request Hits

seqreadreqs
Specifies Search/Read Sequential I/O Requests

seqreadhits
Specifies Search/Read Sequential I/O Request Hits

seqwritereq
Specifies Write Sequential I/O Requests

seqwritehits
Specifies DASD Fast Write Sequential I/O Request Hits

cachfwrreqs
Specifies Search/Read Cache Fast Write I/O Requests

cachfwrhits
Specifies Search/Read Cache Fast Write I/O Request Hits

cachfwreqs
Specifies Cache Fast Write I/O Requests

cachfwhits
Specifies Cache Fast Write I/O Requests Hits

inbcachload
Specifies Inhibit Cache Loading I/O Requests that operate with DASD

bypasscach
Specifies Bypass Cache I/O Requests

seqDASDtrans
Specifies Sequential DASD to Cache Transfer Operations

DASDtrans
Specifies DASD to Cache Transfer Operation Count

cachetrans
Specifies Cache to DASD Transfer Operation Count

NVSpadel
Specifies DASD Fast Write Operations Delayed Due to nonvolatile storage Space Constraints

normwriteops
Specifies Normal 'DASD Fast Write' Write Operation Counts

seqwriteops
Specifies Sequential Access 'DASD Fast Write' Write Operation Counts

reccachemis
Specifies Number of record cache Read Misses

qwriteprotos
Specifies Quick Write Promotes

CKDirtrkac
Specifies Irregular Track Accesses

CKDirtrkhits
Specifies Irregular Track Accesses Hits

cachspdelay

Specifies Operations Delayed Due To Cache Space Constraints

timelowifact

Specifies Milliseconds of lower interface I/O activity for the indicated device.

phread

Specifies Physical Storage Read Operations

phwrite

Specifies Physical Storage Write Operations

phbyteread

Specifies Physical Storage Bytes Read in 128 KB increments.

phbytewrit

Specifies Physical Storage Bytes Written in 128 KB increments.

recmoreads

Specifies Record Mode Read Operations

sfiletrkreads

Specifies the Number of tracks read from the Concurrent Copy or XRC Sidefile.

contamwrts

Specifies the Number of Contaminating writes for a Concurrent Copy or XRC volume

PPRCtrks

Specifies the Number of tracks or portion of tracks that were transferred to the secondary device of a PPRC pair.

NVSpallo

Specifies the NVS Space Allocations

timephread

Specifies the physical storage read response time in 16 ms increments.

timephwrite

Specifies the physical storage write response time in 16 ms increments.

byteread

Specifies the number of bytes read in 128 KB increments

bytewrit

Specifies the number of bytes written in 128 KB increments.

timeread

Specifies the accumulated response time for all read operations.

timewrite

Specifies the accumulated response time for all write operations.

Logical subsystem specific commands

This section contains commands that are used to modify logical subsystems for Open Systems and to display logical subsystem information.

Use the following commands to modify logical subsystems and to display logical subsystem information:

- **chlss**
- **lslss**
- **showlss**

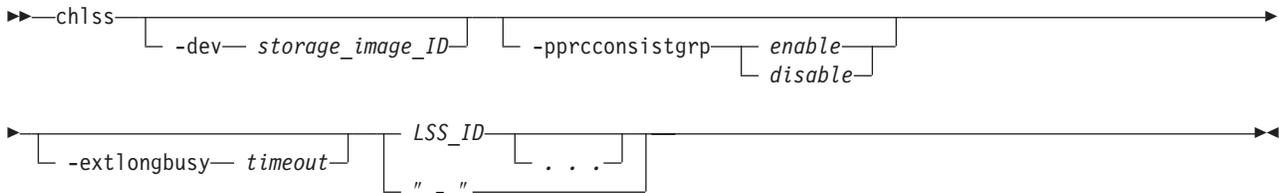
The **chlss** command modifies one or more logical subsystems.

The **lslls** command 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.

The **showlls** command generates a report that displays the detailed properties of a specified LSS.

chlss

The **chlss** command modifies a logical subsystem.



Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for all logical subsystems (LSSs); otherwise, it is optional.

-pprconsistgrp *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 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.

LSS_ID . . . | -

(Required) Specifies one or more LSSs to be modified by this command. An LSS ID is two hexadecimal characters 00 - FE (2107 only) .

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.

If you specify a dash (-), this parameter information is automatically supplied.

Example: 00-03 08

Example (2107)

Invoking the chlss command

```
dslcli>chlss -dev IBM.2107-75FA120 06 0F
```

The resulting output


```
dsccli>ls|ss -dev IBM.1750-75FA120 -1
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

ID	Group	Addrgrp	Stgtype	Confgvols	Subsys
IBM.2107-75FA120/10	0	1	fb	256	FF10
IBM.2107-75FA120/11	1	1	fb	256	FF11
IBM.2107-75FA120/12	0	1	fb	256	FF12

Report field definitions

ID Specifies 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 - FF (2107 only).

Group Specifies the server that is managing the logical subsystem group. The server is identified as either 0 or 1.

Addrgrp

Specifies the address group object of which the logical subsystem is a member.

Stgtype

Specifies 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

Specifies the number of volumes currently assigned to this logical subsystem.

Subsys

Specifies the user assigned, or default SSID value.

showlss

The **showlss** command displays detailed properties of a logical subsystem (LSS).

```
►► showlss [ -dev storage_image_ID ] [ LSS_ID ]
```

Parameters

-dev storage_image_ID

(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. This parameter is required if you do not specify a fully qualified ID for the logical subsystem; otherwise, it is optional.

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 (2107 only).

The following is an example of a fully qualified LSS ID: IBM.2107-75FA120/10

The following is an example of a shortened version of the LSS ID when the **dev** parameter is specified:

```
dscli>showlss -dev IBM.2107-75FA120 10
```

If you specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the showlss command to show default information

```
dscli>showlss IBM.2107-75FA120/10
```

The resulting output

```
DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120  
Sun Aug 11 02:23:49 PST 2004 IBM
```

ID	Group	Addrgrp	Stgtype	Confgvols
IBM.2107-75FA120/10	0	1	fb	256

Pprconsistgrp	Xtndlbztimeout (secs)	Subsys
Disabled	120	FF10

Report field definitions

ID Specifies the unique identifier that is assigned to this logical subsystem. It includes the storage image ID and the 2 digit LSS ID 00 - FE (2107 only) .

Group Specifies the server that manages the logical subsystem. The displayed values are 0 or 1.

Addrgrp

Specifies the address group object that the logical subsystem is a member of.

Confgvols

Specifies the number of volumes that are assigned to this logical subsystem.

Stgtype

Specifies the type of storage volumes contained by this logical subsystem. The value displayed is fb (fixed block) or ckd (count key data)

Pprconsistgrp

Specifies the assigned PPRC consistency group state setting. If you do not designate enabled, the default value of disabled is assigned.

Xtndlbztimeout (secs)

Specifies the assigned extended long busy timeout value. If you do not designate a value, the default value of 210 seconds is assigned and displayed.

Subsys

Specifies the user assigned, or default SSID value.

Fixed block logical volume specific commands

This section contains commands that are used to create, modify, and delete fixed block logical volumes for Open Systems and to display logical volume information.

Use the following commands to create, modify, and delete fixed block logical volumes and to display logical volume information:

- **chfbvol**
- **lsfbvol**
- **mkfbvol**
- **rmfbvol**
- **showfbvol**

The **chfbvol** command modifies the name or data type of a fixed block volume.

The **lsfbvol** command 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.

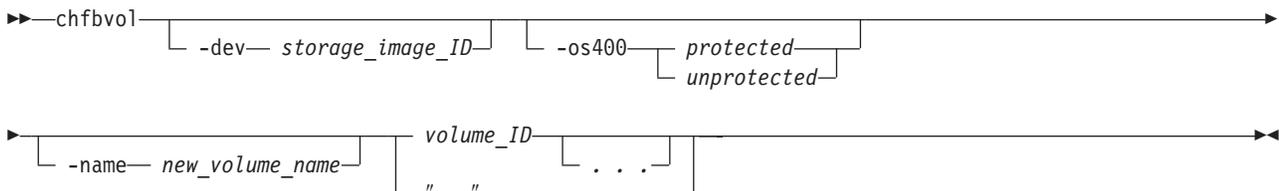
The **mkfbvol** command creates open systems fixed block volumes in a storage unit.

The **rmfbvol** command deletes one or more specified fixed block volumes from a storage unit.

The **showfbvol** command 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 changes the name or data type 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. This parameter is required if you do not specify a fully qualified volume ID.

-os400 *protected | unprotected*

(Optional) Specifies options for the OS 400 volume data types. If the current volume data type is os400-protected(FB 520P) or os400-unprotected(FB 520U) then this data type can be modified to either os400-unprotected(FB 520U) or os400-protected(FB520P) respectively. If the current volume data type is 512, then this setting cannot be modified.

-name *new_volume_name*

(Optional) Specifies your nickname for this volume. A nickname cannot exceed 16 characters.

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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF). You must fully qualify the volume ID with manufacturer, 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

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the chfbvol command

```
dscli>chfbvol -dev IBM.2107-75FA120 -os400 protected 0100 0101
```

The resulting output

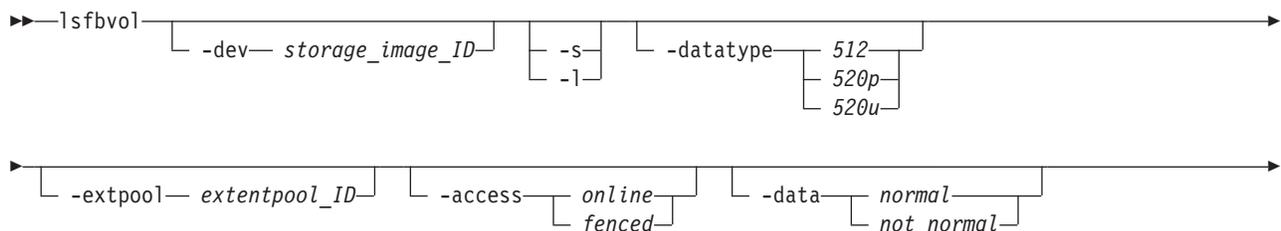
```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

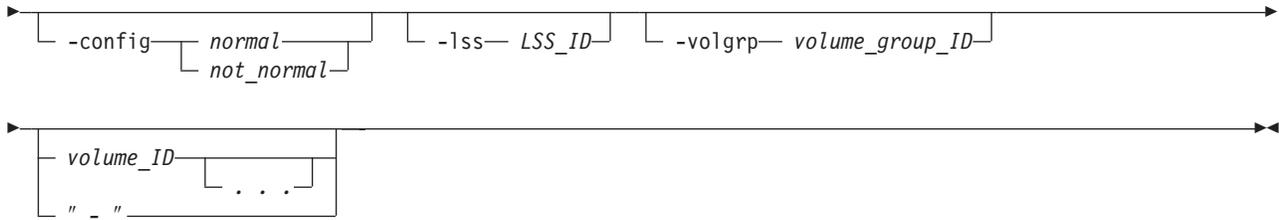
```
FB volume 0100 successfully modified.
```

```
FB volume 0101 successfully modified.
```

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, it is recommended that you specify the **-lss** or the **-volgrp** parameter as part of your command.

- dev** *storage_image_ID*
(Optional) Specifies the storage image ID, which consists of manufacturer, type, and serial number.
- s** (Optional) Displays the volume IDs. You cannot use the **-l** and the **-s** parameters together.
- l** (Optional) Displays default output plus additional attributes that are identified as long output. You cannot use the **-l** and the **-s** parameters together.
- datatype** *512 | 520p | 520u*
(Optional) Displays volumes of the specified volume data type. Standard 2107/1750 volume (512), System i protected (520p), System i unprotected (520u).
- 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 2107).
- 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.
- 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 X is the address group (0 - 1), XY together is the logical subsystem number 00 - FE, and ZZ is 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.

If you specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both machine types. The only difference is the machine type designation, 2107 versus 1750.

Note:

Invoking the lsfbvol command

```
dscli>lsfbvol
-dev IBM.2107-75FA120 -l -volgrp V2
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120

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

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	-

Report field definitions

Name Identifies the nickname that you assigned for the specified volume object.

ID Identifies 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

Specifies 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.

Indeterminate data loss

Specifies 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

Specifies 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

Specifies 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

Specifies 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.

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.

Migrating

Indicates that the logical volume is in the process of performing dynamic volume relocation to a specified extent pool.

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.

Corrective action: Use the **rmfbvol** command to delete each volume 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.

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.

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-A0x

Indicates that the *xxxx* is a 2107 or 1750. The A0 indicates a System i protected volume (for example, 2107-A01 or 1750-A07).

xxxx-A8x

Indicates that 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 binary (B) values will be displayed:

- FB 512
- FB 520P
- FB 520U

Extpool

Identifies the extent pool ID. Volume extents are allocated from this extent pool ID.

SAM (Storage Allocation Method)

Specifies the space efficient implementation. The current designation is Standard.

Capttype

Indicates the capacity unit type that is used at volume creation. One of the following values is displayed:

ESS

The capacity unit is 10^9 B.

DS

The capacity unit is 2^{30} B.

DS/ESS

The capacity unit is 2^{30} B or 10^9 B.

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^{30} B)

Specifies the size of the volume that is available for host system access in 2^{30} B (binary GB) units.

Note: "-" (null) is displayed if the capacity unit type of the volume is ESS (capttype=ESS)

Cap (10^9 B)

Specifies the size of the volume that is available for host system access in 10^9 B (decimal GB) units.

Note: "-" (null) 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

Specifies 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 null (-) is displayed if there are no volume groups that are associated with the volume.

mkfbvol

The **mkfbvol** command creates open systems fixed block (FB) volumes in a storage image.



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.

-dev storage_image_ID

(Optional) Specifies the storage image ID, which includes manufacturer, type, and serial number.

-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 zeroes, prefixed with the letter *P*.

-os400 *A01 | A81 | A02 | A82 | A04 | A84 | A05 | A85 | A06 | A86 | A07 | A87*

(Optional) The OS 400 volume options. If this flag is not specified, the default standard 2107/1750 volume is created.

This parameter is required if capacity is not specified.

The storage sizes and the data types for this volume:

- A01 - 8.6 GB, protected
- A81 - 8.6 GB, unprotected
- A02 - 17.5 GB, protected
- A82 - 17.5 GB, unprotected
- A04 - 70.5 GB, protected
- A84 - 70.5 GB, unprotected
- A05 - 35.1 GB, protected
- A85 - 35.1 GB, unprotected
- A06 - 141.12 GB, protected
- A86 - 141.12 GB, unprotected
- A07 - 282.25 GB, protected
- A87 - 282.25 GB, unprotected

GB is 10^9 Byte.

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 10^9 Bytes.

ds: Specifies that the unit is 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.

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:

- 1GB = 1,073,741,824 (2^{30} Bytes)
- Supported storage sizes range from 1 to 2048.

If the **-type ess** parameter is specified:

- capacity = X.Y or X
 - where X is whole gigabytes, with 1 GB=1,000,000,000 (10⁹ Bytes).
 - where Y represents a fraction of 1 GB. Y is limited to a single digit (0...9) to the right of the decimal.
- Supported storage sizes range from 0.1 to 982.2 (0.1 increment).

If the *-type blocks* parameter is specified, *capacity* is the number of 512 blocks. Supported storage sizes are from 1 to 4294967296 (4x2³⁰ Bytes).

-name *volume_name*

(Optional) Your nickname for this volume. The nickname can be 16 characters long, and it 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 zeroes, 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.

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 is a 32 bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is 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

If you specify a dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the mkfbvol command

```
dsccli>mkfbvol -dev IBM.2107-75FA120
-extpool P1 -name my_vol_#d -type ess
-cap 8.6 0100 0101 0102 0103
```

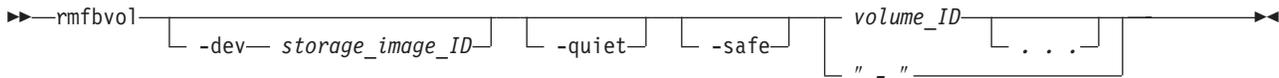
The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Device: IBM.2107-75FA120
```

```
FB volume 0100 successfully created.
FB volume 0101 successfully created.
FB volume 0102 successfully created.
FB volume 0103 successfully created.
```

rmfbvol

The **rmfbvol** command deletes fixed block volumes from a storage image.



Parameters

Note: 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.

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. This parameter is required if you do not specify a fully qualified volume ID for all logical volumes and you do not specify a value for the *dev* variable in your profile file.

-quiet

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

-safe

(Optional) Specifies that the system perform a check to see if the specified volumes are assigned to any user-defined volume group. If any volumes are still assigned to a user-defined volume group, the **rmfbvol** command fails without deleting any volumes. When this occurs, messages are provided that list the volumes that are still assigned to a user-defined volume group.

Note: If there is any reason that the system cannot perform the check, the **rmfbvol** command fails and no volumes are deleted.

volume_ID . . . | -

(Required) Specifies an array of one or more fixed block volume IDs to be removed. This parameter also accepts a fully qualified volume ID, which includes the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified.

Example of **-dev** parameter use

If you specify the **-dev** parameter, you can use the shortened version of the *Volume_ID* parameter as follows:

```
dscli>rmfbvol -dev IBM.2107-75FA120 0100-010F 0180-018F 0120
```

If you do not specify the **-dev** parameter and you specify the *volume_ID* parameter, you must use the fully qualified version of the volume ID as follows:

```
dscli>rmfbvol IBM.2107-75FA120/0100-010F 0180-018F 0120
```

The volume ID is a 32 bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is 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.

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the rmfbvol command

```
dscli>rmfbvol -dev IBM.2107-75FA120 0100 0101
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
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 is an example of the output that results when you specify that you want a range of volume IDs to be removed.

Invoking the rmfbvol command

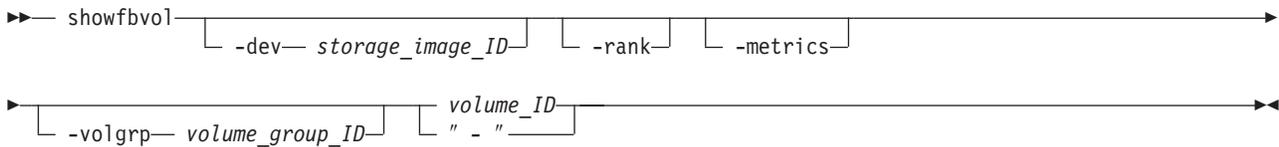
```
dscli>rmfbvol -dev IBM.2107-75FA120 0005-00FF
```

```
Date/Time: September 13, 2005 9:38:52 PM MST IBM DSCLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
CMUC00027W rmfbvol: Are you sure you want to delete FB volume 0005-00FF?  
[y/n]:y  
CMUC00028I rmfbvol: FB volume 0005 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0006 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0007 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0008 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0009 successfully deleted.  
CMUC00028I rmfbvol: FB volume 000A successfully deleted.  
CMUC00028I rmfbvol: FB volume 000B successfully deleted.  
CMUC00028I rmfbvol: FB volume 000C successfully deleted.  
CMUC00028I rmfbvol: FB volume 000D successfully deleted.  
CMUC00028I rmfbvol: FB volume 000E successfully deleted.  
CMUC00028I rmfbvol: FB volume 000F successfully deleted.  
CMUC00028I rmfbvol: FB volume 0010 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0011 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0012 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0013 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0014 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0015 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0016 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0017 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0018 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0019 successfully deleted.  
CMUC00028I rmfbvol: FB volume 001A successfully deleted.  
CMUC00028I rmfbvol: FB volume 001B successfully deleted.  
CMUC00028I rmfbvol: FB volume 001C successfully deleted.  
CMUC00028I rmfbvol: FB volume 001D successfully deleted.  
CMUC00028I rmfbvol: FB volume 001E successfully deleted.  
CMUC00028I rmfbvol: FB volume 001F successfully deleted.  
CMUC00028I rmfbvol: FB volume 0020 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0021 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0022 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0023 successfully deleted.  
CMUC00028I rmfbvol: FB volume 0024 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 the manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the logical volume.

-rank

(Optional) Specifies that a rank extents table 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: This parameter cannot be used with the **-metrics** or **-volgrp** parameters.

-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 **-rank** 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 be displayed.

Notes:

1. You can only use the **-volgrp** parameter 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 **-rank** or **-metrics** 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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE, and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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.

A separate example is not shown for the 1750 because the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

Invoking the showfbvol to show volume properties

Note: The example output is based on using the **showfbvol** command for a 1.0 (Binary) GB volume.

```
dscli>showfbvol  
-dev IBM.2107-1300861 -rank 6000
```

The resulting output

Date/Time: April 13, 2005 2:38:14 PM JST IBM DS CLI Version: 5.0.3.40 DS: IBM.2107-1300861

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
P0	1	DS	1.0	-	2097152	V2	3

Rank	Extents
R0	1
R2	2

Report field definitions (*-metrics* parameter not specified)

Name

Specifies the nickname that you assigned for this volume object.

ID Specifies 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.

Datstate

One of the following data states are displayed:

Normal

None of the other data states apply. The access state is Online.

Pinned

Specifies 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

Specifies 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. This data state transitions to Rank repairing 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

Specifies 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

Specifies 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

Specifies 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.

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.

Migrating

Indicates that the logical volume is in the process of performing dynamic volume relocation to a specified extent pool.

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.

Corrective action: Use the **rmfbvol** command to delete each volume 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.

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.

deviceMTM

Indicates the volume device type and the model. The volume MTM is determined by the fixed block volume data type and the volume capacity (in GB). The model is either 2107 or 1750 however, the MTM can be any one of the following depending on your system:

2107-900

This indicates a standard 2107 volume.

1750-500

This indicates a standard 1750 volume.

xxxx-A0x

The *xxxx* is 2107 or 1750; the A0 indicates an System i protected volume (for example, 2107-A01 or 1750-A07).

xxxx-A8x

The *xxxx* is 2107 or 1750; the A8 indicates an System i unprotected volume (for example, 2107-A81 or 1750-A87).

Datatype

Indicates the volume data type setting. One of the following values will be displayed:

- FB 512
- FB 520P
- FB 520U

Addrgrp

Specifies the address group that contains the designated volume object. An address group ID is one hexadecimal character (0 - F).

Extpool

Specifies the extent pool ID. Volume extents are allocated from this extent pool ID.

Exts

Specifies the number of extents used by the designated volume ID.

Capttype

Indicates capacity unit type used at volume creation. One of the following values is displayed:

ESS

The capacity unit is 10^9 B.

DS

The capacity unit is 2^{30} B.

DS/ESS

The capacity unit is 2^{30} B or 10^9 B.

Blocks

The capacity unit 512B.

iSeries

The capacity unit was not specified at volume creation. This fixed block volume was created for iSeries.

Cap (2^{30} B)

Specifies the size of volume that is available for host system access in 2^{30} B (binary GB) unit.

Note: "-" (null) is displayed if the capacity unit type of the volume is ESS (capttype=ESS)

Cap (10^9 B)

Specifies the size of volume that is available for host system access in 10^9 B (decimal GB) unit.

Note: "-" (null) 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

Specifies 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 null (-) is displayed if there are no volume groups that are associated with the volume.

Ranks

Specifies the number of ranks the volume resides on.

Rank (Rank Extent table)

Specifies the rank ID.

Extents (Rank Extents table)

Specifies 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

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI
Version: 5.0.0.0 DS: IBM.2107-75FA120
```

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	cachfwr-reqs	cachfwr-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	time-phread	timeph-write	byte-read
10000	10000	10000	10000	10000	10000	10000	10000

bytewrit	timeread	timewrite
10000	10000	10000

Report field definitions (*-metrics* parameter specified)

ID Specifies the unique identifier that is assigned to this volume object.

Date

Specifies the current time stamp for the volume performance counters.

normrdqrts

Specifies Search/Read Normal I/O Requests.

normrdhits

Specifies Search/Read Normal I/O Requests instances.

normwritereq

Specifies Write Normal I/O Requests.

normwritehits

Specifies DASD Fast Write I/O Request instances.

seqreadreqs

Specifies Search/Read Sequential I/O Requests.

seqreadhits

Specifies Search/Read Sequential I/O Request instances.

seqwritereq

Specifies Write Sequential I/O Requests.

seqwritehits

Specifies DASD Fast Write Sequential I/O Request instances.

cachfwrreqs

Specifies Search/Read Cache Fast Write I/O Requests.

cachfwrhits

Specifies Search/Read Cache Fast Write I/O Request instances.

cachfwreqs

Specifies Cache Fast Write I/O Requests.

cachfwhits
Specifies Cache Fast Write I/O Requests instances.

inbcachload
Specifies Inhibit Cache Loading I/O Requests that operate with DASD.

bypasscach
Specifies Bypass Cache I/O Requests.

seqDASDtrans
Specifies Sequential DASD to Cache Transfer Operations.

DASDtrans
Specifies DASD to Cache Transfer Operation Count.

cachetrans
Specifies Cache to DASD Transfer Operation Count.

NVSpadel
Specifies DASD Fast Write Operations Delayed Due to nonvolatile storage Space Constraints.

normwriteops
Specifies Normal 'DASD Fast Write' Write Operation Counts.

seqwriteops
Specifies Sequential Access 'DASD Fast Write' Write Operation Counts.

reccachemis
Specifies Number of record cache Read Misses.

qwriteprots
Specifies Quick Write Promotes.

CKDirtrak
Specifies Irregular Track Accesses.

CKDirtrkhits
Specifies Irregular Track Accesses instances.

cachspdlay
Specifies Operations Delayed Due To Cache Space Constraints.

timelowifact
Specifies Milliseconds of lower interface I/O activity for the indicated device.

phread
Specifies Physical Storage Read Operations.

phwrite
Specifies Physical Storage Write Operations.

phbyteread
Specifies Physical Storage Bytes Read in 128 KB increments.

phbytewrit
Specifies Physical Storage Bytes Written in 128 KB increments.

recmoreads
Specifies Record Mode Read Operations.

sfiletrkreads
Specifies the Number of tracks read from the Concurrent Copy or XRC Sidefile.

contamwrts

Specifies the Number of Contaminating writes for a Concurrent Copy or XRC volume.

PPRCtrks

Specifies the Number of tracks or portion of tracks that were transferred to the secondary device of a PPRC pair.

NVSpallo

Specifies the NVS Space Allocations.

timephread

Specifies the Physical Storage Read Response Time in 16 ms increments.

timephwrite

Specifies the Physical Storage Write Response Time in 16 ms increments.

byteread

Specifies the number of Bytes read in 128 KB increments.

bytewrit

Specifies the number of Bytes written in 128 KB increments.

timeread

Specifies the accumulated response time for all read operations.

timewrite

Specifies the accumulated response time for all write operations.

Volume group specific commands

This section contains commands that are used to create, modify, and delete volume groups and to display volume group information.

Use the following commands to create, modify, and delete volume groups and to display volume group information:

- **chvolgrp**
- **lsvolgrp**
- **mkvolgrp**
- **rmvolgrp**
- **showvolgrp**

The **chvolgrp** command modifies a volume group name and volume members.

The **lsvolgrp** command 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.

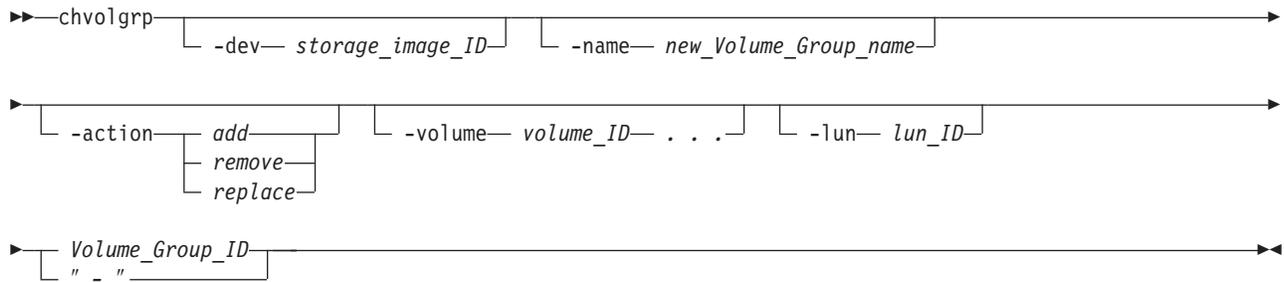
The **mkvolgrp** command creates a volume group in a storage unit.

The **rmvolgrp** command deletes the specified volume group or volume groups from a storage unit.

The **showvolgrp** command 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 systems, 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 **mapscsi256**.

-dev *storage_image_ID*

(Optional). Specifies the storage image ID, which consists of manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the volume group ID.

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.

-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 have specified the **-volume** parameter and not included the **-action** parameter.

-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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

To specify a range of volume IDs, separate the volume IDs with a dash (-). You must separate multiple volume IDs or ranges of volume IDs with a comma between each ID or range of IDs.

Notes:

1. For SCSI MAP 256, the array or ranges cannot exceed 256 volume ID entries. Otherwise, up to 64 384 entries are allowed.
2. The **chvolgrp** command fails if you specify the **-volume** parameter and do not specify the **-action** parameter.

Example: 0100-010F,0180-018E,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 works:

A volume group of "SCSI Map 256" type has Volume 0000 and 0001. Their LUNs are the members of the following volume group:
(showvolgrp displays the current mapping.)

0000 : 10

0001 : 11

Because the range of LUN IDs is 00-FF, the unused LUN IDs are 00,01,...,0F,12,13,...,FF.

If you add volume 0002 and 0003 to this volume group without the **-lun** parameter, the mapping results in the following because 00 and 01 are "smaller" unused LUN IDs:

0002 : 00

0003 : 01

0000 : 10

0001 : 11

If the **-action replace** parameter is specified without specifying the **-lun** parameter, **lun_ID=00** is assumed.

Volume_Group_ID | -

(Required). Specifies the ID of the volume group 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 zeroes, prefixed with the letter V.

If you specify the dash (-), this parameter information is automatically supplied.

Example: IBM.2107-75FA1243/V1231

Example (2107)

Invoking the chvolgrp command

```
dscli>chvolgrp -action add -volume 0000-000F IBM.2107-75FA120/V2341
```

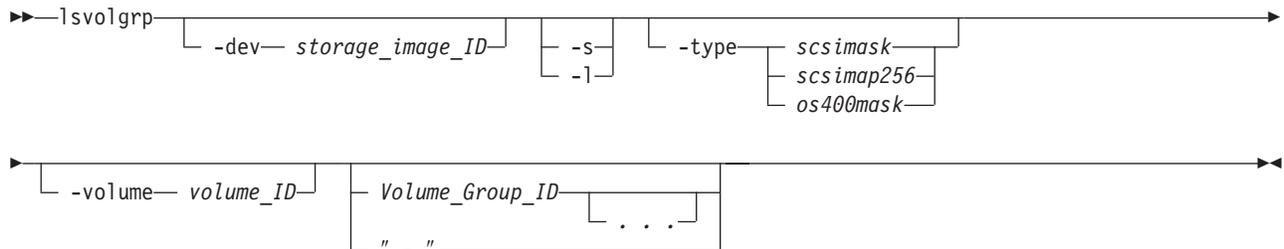
The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
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. This parameter is required if you do not specify a fully qualified volume group ID and you do not specify a value for the *devid* variable in your profile file.

Example: IBM.2107-75FA120

-s (Optional). Displays volume group IDs only. You cannot use the **-l** and the **-s** parameters together.

-l (Optional). Displays the default output plus users. You cannot use the **-l** and the **-s** parameters together.

-type | *scsimask* | *scsimap256* | *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* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is 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.

If you specify a dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both the 2107 and 1750.

Invoking the lsvolgrp command

```
dsccli>lsvolgrp -dev IBM.2107-75FA120 -l
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Device: IBM.2107-75FA120

Name	ID	Type
Host_xxx _volumes	1011	OS400 Mask
Host_yyy _volumes	1111	OS400 Map 256
Host_zzz _volumes	1211	SCSI Mask

Report field definitions

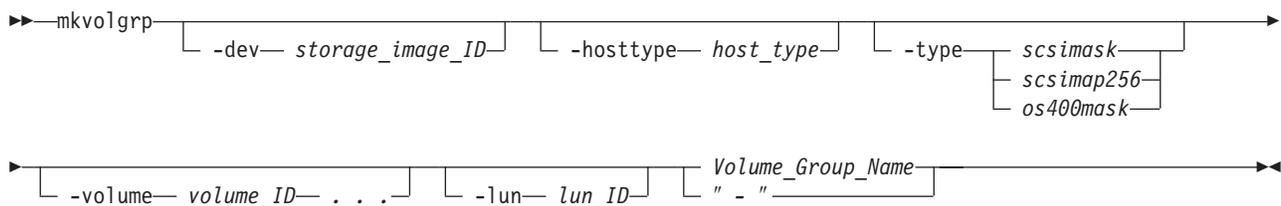
Name Specifies the name you assigned for this volume group ID.

ID Specifies the storage unit ID followed by the volume group ID. The volume group identifier is a four-digit decimal number, with no leading zeros, prefixed by the letter V.

Type Specifies the type of the volume group.

mkvolgrp

The **mkvolgrp** command creates a volume group in a storage image.



Parameters

Notes:

1. It is highly recommended that you use the **-hosttype** parameter when you issue this command.

If you are using an HP-UX operating system, the number of volumes in the volume group must not exceed 7 volumes. This restriction only applies when the **hostconnect** attribute for the **-addrdiscovery** parameter is set to **reportlun** and the associated volume group is of type **scsimap256**.

-dev *storage_image_ID*

(Optional). Specifies the storage image ID, which includes manufacturer, machine type, and serial number. This parameter is required if you do not specify a fully qualified volume group name.

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 OS400 mask volume group. The IBM iSeries 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 OS400 MASK, 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 is the address group (0 - 1), XY together is the logical subsystem number 00 - FE (for 2107), and ZZ is 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 specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the **mkvolgrp** command

```
dscli>mkvolgrp -dev IBM.2107-75FA120
-volume 0000-000F host_xyz_volumes
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
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, type, and serial number. This parameter is required if you do not specify a fully qualified volume group ID and you do not specify a value for the *devid* variable in your profile file.

Example: IBM.2107-75FA120

-quiet

(Optional). Specifies that confirmation prompt for this command be turned off.

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.

If you use the dash (-), this parameter information is automatically supplied.

Example of -dev parameter use

If you specify the **-dev** parameter, you can use the shortened version of the **Volume_Group_ID** parameter as follows:

```
dscli>rmvolgrp -dev IBM.2107-75FA120 V11
```

If you do not specify the **-dev** parameter and you specify the **Volume_Group_ID** parameter, you must use the fully qualified version of the volume group ID as follows:

```
dscli>rmvolgrp IBM.2107-75FA120/V11
```

Example (2107)

Invoking the rmvolgrp command

```
dscli>rmvolgrp IBM.2107-75FA1243/V123
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

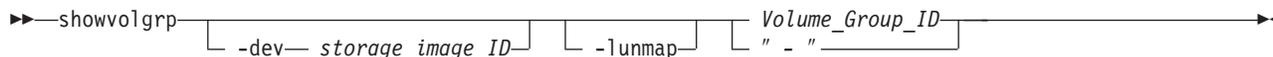
```
Are you sure you want to delete Volume Group IBM.2107-75FA1243/V123? y/n
```

```
Y
```

```
Volume group IBM.2107-75FA1243/V123 successfully deleted.
```

showvolgrp

The **showvolgrp** command displays detailed properties of a volume group.



Parameters

-dev *storage_image_ID*

(Optional). Specifies the storage image ID, which consists of manufacturer, machine type, and serial number.

Note: This parameter is required if you do not specify a fully qualified volume group ID and you do not specify a value for the *dev* variable in your profile file.

Example: IBM.2107-75FA120

-lunmap

(Optional). Specifies that a LUN mapping table be displayed that shows the volume ID and LUN ID relationship. This parameter is valid for all scsi and os400 type volume groups.

Volume_Group_ID | -

(Required). Specifies that the properties be displayed for the specified volume group. This parameter accepts a fully qualified volume group ID, which consists of the storage image ID, or a shortened version without the storage image ID if the **-dev** parameter is specified. The shortened version is a four-digit decimal number with no leading zeros, prefixed with the letter V.

Examples of -dev parameter use

If you specify the **-dev** parameter, you can use the shortened version of the **Volume_Group_ID** parameter as follows:

```
dscli>showvolgrp -dev IBM.2107-75FA120 V11
```

where *V11* represents value for the volume group ID.

If you do not specify the **-dev** parameter, and you specify the **Volume_Group_ID** parameter, you must specify the fully qualified version of the **Volume_Group_ID** parameter as follows:

```
dscli>showvolgrp IBM.2107-75FA120/V11
```

If you specify the dash (-), this parameter information is automatically supplied.

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 **showvolgrp** command. A separate example is not shown for the 1750 because the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

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

Date/Time: May 17, 2006 8:33:58 AM GMT IBM DS CLI Version: 5.0.0.0

DS: IBM.2107-1300861

Name	ID	Type	Vols
My_host_ system_ volumes	V2	SCSI- MAP 256	1000 1001 1002 1003 1004 1005 1006 1007

=====LUN Mapping=====

Vol	LUN
1000	00
1001	01
1002	02
1003	03
1004	04
1005	05
1006	06
1007	07

Invoking the showvolgrp command where the volume group type is SCSI Mask
 dscli>showvolgrp -lunmap IBM.2107-1300861/V18

The resulting output

Date/Time: May 17, 2006 8:38:08 AM GMT IBM DS CLI Version: 5.0.0.0
 DS: IBM.2107-1300861

Name	ID	Type	Vols
myVG1	V18	SCSI-Mask	1000 1001 1002 1003 1004 1005 1006 1007

=====LUN Mapping=====

Vol	LUN
1000	40104000
1001	40104001
1002	40104002
1003	40104003
1004	40104004
1005	40104005
1006	40104006
1007	40104007

Report field definitions

- Name** Specifies the name that you assigned for the designated volume group ID.
- ID** Specifies the volume group ID. The volume group identifier is a four-digit decimal number having no leading zeros, and prefixed by a V.
- Type** Specifies the configured volume group type. Any one of the following volume group types can be queried: SCSI all | SCSI Mask | SCSI MAP 256 | os400 all | os400 Mask

Note: os400 all and os400 Mask are sometimes referred to as SCSI520 all and SCSI520 Mask.

Vols Identifies the complement of accessible volume numbers within the designated volume group.

Vol (part of LUN mapping table)
 Specifies the volume ID.

LUNS (part of LUN mapping table)
 Specifies 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.

Advanced operation commands

This section contains commands that are used to further administer and tune storage.

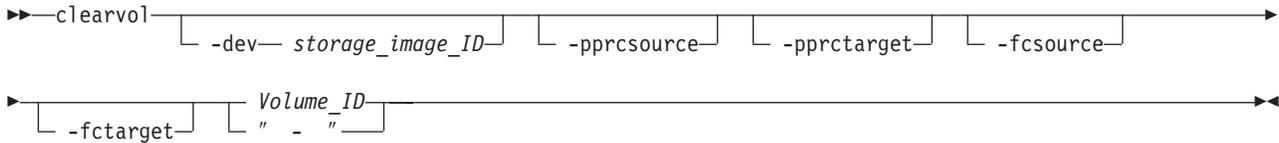
Use the following command to administer and tune storage:

clearvol

The clearvol command clears Copy Services relationships for a base logical volume.

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, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the volume ID.

Example: IBM.2107-75FA120

-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* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

Example: IBM.2107-75FA120/0001

Example (2107)

Invoking the clearvol command

```
dscli>clearvol -dev IBM.2107-75FA120 0001
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
Volume 0001 successfully cleared.
```

Copy Services commands

This section contains Copy Services commands.

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

FlashCopy commands

This section contains commands that are used to configure FlashCopy relationships and to display FlashCopy information.

Use the following commands to configure FlashCopy relationships and to display FlashCopy information:

- **commitflash**
- **resyncflash**
- **lsflash**
- **mkflash**
- **reverseflash**
- **revertflash**
- **rmflash**
- **unfreezeflash**
- **setflashrevert**

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

The **resyncflash** command 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.

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

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

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

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

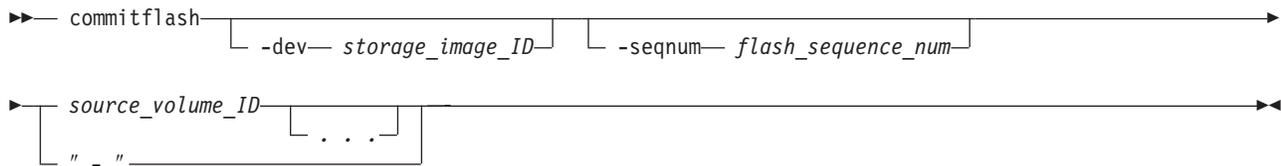
The **rmflash** command removes a relationship between FlashCopy volume pairs.

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

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

commitflash

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



Parameters

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

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

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. This parameter is required if you do not specify a fully qualified ID for all the source volumes and you do not specify a value for the *devid* variable in your profile file.

Example: IBM.2107-75FA120

-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.

pair. Use this parameter if you intend to use the **resyncflash** command again with this pair. The **-persist** parameter is automatically selected when the **-record** parameter is selected.

Note: You cannot reissue the **resyncflash** command if you leave this parameter out of your command string.

-persist

(Optional) Retains the FlashCopy relationship after the background copy completes. The FlashCopy relationship between the source and the target volumes remains indefinitely until broken by a **rmflash** command. This parameter is automatically selected when the **-record** parameter is selected. Select this parameter along with the **-record** parameter if you are using this volume pair with the commands **resyncflash**, **reverseflash**, or **setflashrevertible**.

Note: You cannot reissue the **resyncflash** command if you leave this parameter out of your command string.

-freeze

(Optional) 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*. 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.

-tgtpprc

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

-tgtoffline

(Optional) Causes the **mkflash** command to be rejected if the target volume is online for host system access. This parameter only applies to CKD volumes.

-tgtinhibit

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

-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. Only the relationships that are modified successfully by the **resyncflash** command get the specified sequence number, leaving the ones that failed with the previous one (if they were previously specified).

This parameter is not supported for machine type 2105.

Example: 0010

source_volume_ID:target_volume_ID . . . | -

(Required) Increments a FlashCopy relationship 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 specify a value for the *dev* 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* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the resyncflash command

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

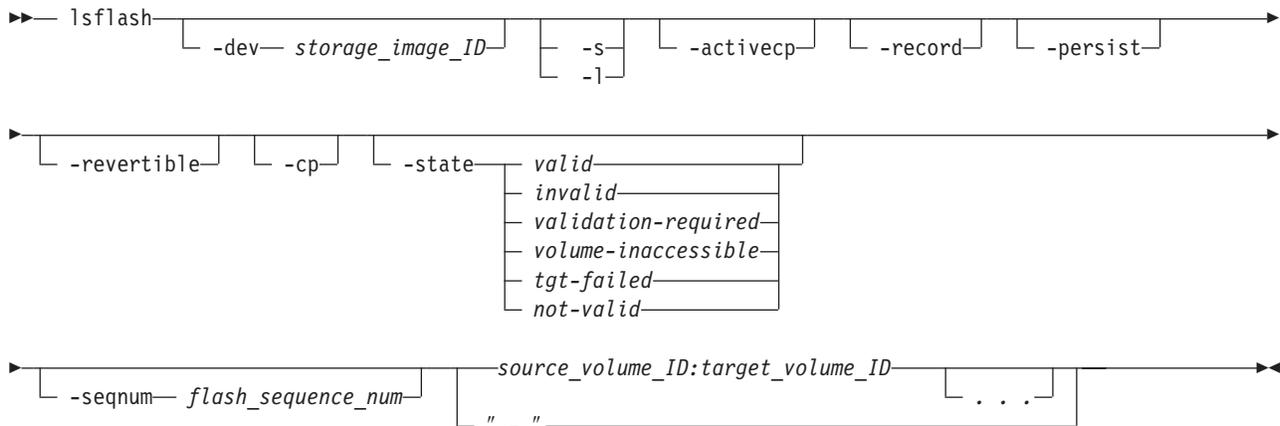
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

```
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, type,

and serial number. 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: IBM.2107-75FA120

- s (Optional) Displays FlashCopy pair IDs. You cannot use the -l and the -s parameters together.
- l (Optional) Displays the default output plus *copy indicator, out of sync tracks, date created, and date synchronized*. You cannot use the -l and the -s 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.

-record

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

-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.

-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 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.

This parameter is not supported for machine type 2105.

source_volume_ID:target_volume_ID | source_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 source volume ID is 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 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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is 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 source volume IDs and a range of source volume IDs. You cannot specify a combination of FlashCopy pair IDs and source volumes IDs.

If you specify the dash (-), this parameter information is automatically supplied.

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

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. A separate example is not shown for the 1750 because the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

Invoking the lsflash command

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

The resulting output

Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0

DS: IBM.2107-75FA120

ID	SrcLSS	Sequence Num	Time-out (secs)	Active Copy	Recording	Persistent	Reversible
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	Background-Copy	Copy-Indicator	OutOf-Sync-Tracks	Date-Created	Date-Synced	State
Enabled	Disabled	Disabled	Yes	0	12/01/2003 02:20:00	12/01/2003 02:23:47	Valid
Enabled	Disabled	Disabled	Yes	0	12/01/2003 02:20:00	12/01/2003 02:23:47	Valid

Source-Write-Enabled	Target-Write-Enabled	Back-ground-Copy	Copy-Indicator	OutOf-Sync-Tracks	Date-Created	Date-Synced	State
Enabled	Disabled	Disabled	Yes	0	12/01/2003 02:20:00	12/01/2003 02:23:47	Valid
Enabled	Disabled	Disabled	Yes	0	12/01/2003 02:20:00	12/01/2003 02:23:47	Valid

Report field definitions

ID Specifies 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.

SrcLSS

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

Sequence Num

Specifies the sequence number that is associated with the FlashCopy relationship.

Note: This item is not supported on the 2105.

Timeout (secs)

Specifies the consistency group Long Busy Timeout setting for the LSS ID that is associated with the source volume of this FlashCopy relationship. You can specify a value in the range of 1 - 600 000. The default value is 120 seconds.

ActiveCopy

Specifies (Enabled or Disabled) whether the background copy process is active for this FlashCopy relationship.

Recording

Specifies (Enabled or Disabled) whether this FlashCopy relationship was established with the record changes option.

Persistent

Specifies (Enabled or Disabled) whether this FlashCopy relationship was established with the persistent option.

Revertible

Specifies (Enabled or Disabled) whether this FlashCopy relationship was established with the revertible option.

SourceWriteEnabled

Specifies (Enabled or Disabled) whether this FlashCopy relationship was established with the allow source writes option.

TargetWriteEnabled

Specifies (Enabled or Disabled) whether this FlashCopy relationship was established with the allow target writes option.

BackgroundCopy

Specifies (Enabled or Disabled) whether this FlashCopy relationship was established with the run background copy option.

CopyIndicator

Specifies (Yes or No) whether the CopyIndicator is set for this FlashCopy relationship.

OutOfSyncTracks

Specifies 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.

DateCreated

Specifies the date and the time that the FlashCopy relationship was established.

DateSynced

Specifies the date and time this FlashCopy relationship was synchronized, or null (-) if the relationship is not synchronized.

State Specifies 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 null (-) 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. Generally, 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

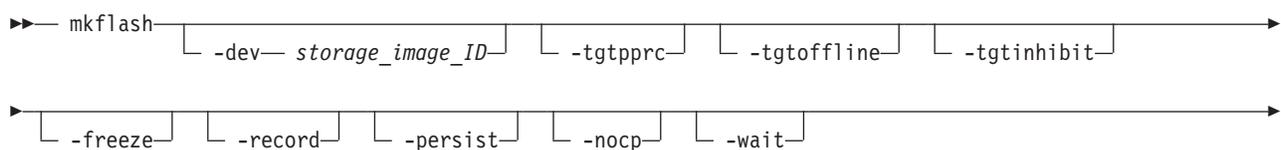
Indicates that the volume cannot be accessed and that the query has failed. When this state is displayed, it generally means that the volume is in a fenced condition.

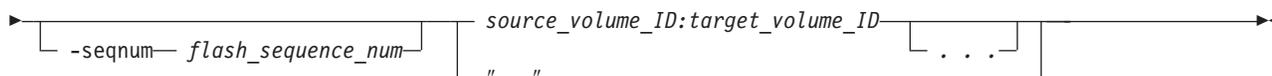
Invalid

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

mkflash

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





Parameters

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which consists of a manufacturer, machine type, and serial number. 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 *dev*id variable in your profile file.

Example: IBM.2107-75FA120

-tgtpprc

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

-tgtoffline

(Optional) Causes the **mkflash** command to be rejected if the target volume is online for host system access. This parameter applies only to count key data volumes.

-tgtinhibit

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

-freeze

(Optional) 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*. 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: If you need to modify the extended long busy timeout setting, use the **chlss** and **chlcu** commands.

-record

(Optional) Records the changed tracks on both volumes within a FlashCopy pair. Select this parameter when you establish an initial FlashCopy volume pair that you intend to use with the **resyncflash** command. The **-persist** parameter is automatically selected when this parameter is selected.

-persist

(Optional) Retains the FlashCopy relationship after the background copy completes. The FlashCopy relationship between the source and the target volumes remains indefinitely until you issue a **rmflash** command. This parameter is automatically selected when the **-record** parameter is selected. Select this parameter along with the **-record** parameter if you will be using this volume pair with the **resyncflash**, **reverseflash**, or **setflashrevertible** commands.

-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.

-wait

(Optional) Delays the command response until the background copy process completes. You cannot use the **-wait** parameter with either the **-persist** or **-nocp** parameters.

-seqnum *flash_sequence_num*

(Optional) Associates the FlashCopy relations that are established with the specified sequence number. This sequence number can be used as an identifier for a relation or group of relations.

This parameter is not supported for machine type 2105.

Example: 0010

source_volume_ID:target_volume_ID . . . | -

(Required) Establishes a FlashCopy relationship for the source and target volume pairs with the specified IDs. This command 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: It is possible that during processing you will receive an error message indicating that relationships have 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. You can issue a **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* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the **mkflash** command

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

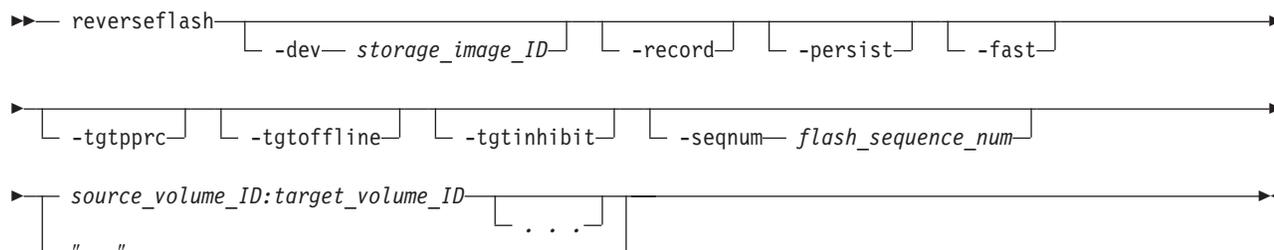
The resulting output

Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

FlashCopy pair 0100:0200 successfully created.

reverseflash

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



Parameters

The direction of a FlashCopy relationship can be reversed, where the volume that was previously defined as the target becomes the source for the volume that was previously defined as the source (and is now the target). 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 volume B is copied to volume A.

The source volume and the target volume that are specified by this command are the current source and target volumes that, as a result of this command process, become reversed. The specified target volume becomes the source and the specified source volume becomes the target. The target parameters specified by this command apply to the new target (the source volume that becomes the target volume when this command is processed).

After the reversal takes place, ensure that you designate this new relationship when you issue any future commands. Failure to designate this reversed relationship could produce unexpected results. For example, you reverse the relationship of source volume 1600 and target volume 1800. Using the **reverseflash** command, your source volume becomes 1800 and your target volume become 1600. All queries and future processing on this relationship must show volume 1800 as your source and volume 1600 as your 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 includes manufacturer, machine type, and serial number. 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 *dev* variable in your profile file.

Example: IBM.2107-75FA120

-record

(Optional) Records the changed tracks on both volumes within a FlashCopy

pair. Use this parameter when you establish an initial FlashCopy volume pair that you intend to use with the **resyncflash** command. The **-persist** parameter is automatically designated when the **-record** parameter is selected.

-persist

(Optional) Retains the FlashCopy relationship after the background copy completes. The FlashCopy relationship between the source and the target volumes remains indefinitely until it is broken by a **rmflash** command. This parameter is automatically designated when the **-record** parameter is specified. Select the **-persist** parameter along with the **-record** parameter if you are using this volume pair with the **resyncflash**, **reverseflash**, or **setflashrevertible** commands.

-fast

(Optional) Specify this parameter when the **reverseflash** command will be issued 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 (processing the **reverseflash** command with the **-fast** parameter), the original FlashCopy target volume is unusable. Normally, after this command completes the background copy, the new FlashCopy target is used as FlashCopy source to restore the original FlashCopy target.

-tgtpprc

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

-tgtoffline

(Optional) Causes the **reverseflash** command to be rejected if the target volume is online for host system access. This parameter only applies to CKD volumes.

-tgtinhibit

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

-seqnum *flash_sequence_num*

(Optional) Associates the FlashCopy relations that are reversed with the specified sequence number. Only the relations that are successfully modified by the command receive the specified sequence number, leaving the ones that failed with the previous number (if previously specified). The sequence number is a four-digit hexadecimal number (0000 - FFFF). The default is 0000.

Example: 0010

This parameter is not supported for machine type 2105.

source_volume_ID:target_volume_ID . . . | -

(Required) Reverses a FlashCopy relationship for the source and target volume pairs with the specified 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 or you specify a value for the *dev* 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* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the reverseflash command

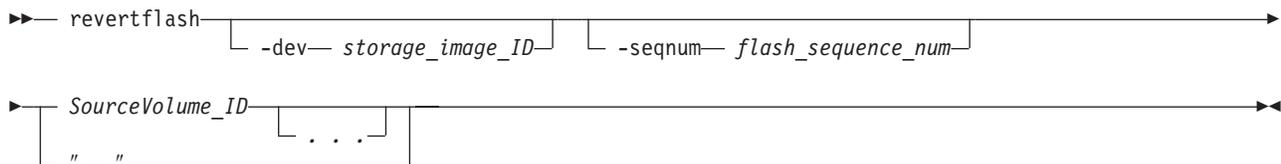
```
dscli>reverseflash -dev IBM.2107-75FA120 0100:0200
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120 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. This parameter is required if you do not specify a fully qualified ID for the source volume and you do not specify a value for the *devid* variable in your profile file.

Example: IBM.2107-75FA120

-seqnum *flash_sequence_num*

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

This parameter is not supported for machine type 2105.

Example: 0010

SourceVolumeID . . . | -

(Required) Specifies the source volume ID for which the FlashCopy relationship is to be reverted. The chosen FlashCopy pair is the one established or modified with the **-record** parameter. This parameter accepts fully qualified volume IDs, which includes storage image IDs, or a shortened version without storage image IDs if the **-dev** parameter is specified or you specify a value for the *dev* 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* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the revertflash command

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

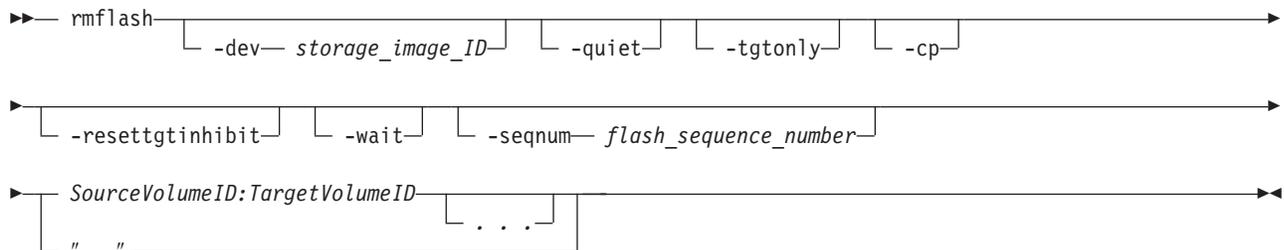
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

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

rmflash

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



Parameters

Notes:

1. Invoking this command with the **-cp** parameter on a FlashCopy relationship that was previously marked with the **-persist** parameter will 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 FlashCopy pairs specified.

4. Do not use the **-wait** parameter on persistent relations.
5. The **-seqnum** parameter is not supported for model 2105.

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, machine type, and serial number. 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 *devuid* variable in your profile file.

Example: IBM.2107-75FA120

-quiet

(Optional) Specifies that the confirmation prompt for this command be turned off.

-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. This parameter applies only to fixed block volumes.

-cp

(Optional) Specifies that the FlashCopy relationship 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 specified in the command.

-resettgtinhibit

(Optional) Specifies that the flag that does not allow host system write operations to the target ID while the FlashCopy relationship exists 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 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 occurs when this is done.

-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 machine type 2105.

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 *dev*id 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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the **rmflash** command

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

The resulting output

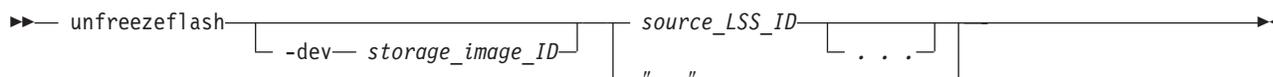
```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
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. This parameter is required if you do not specify a fully qualified source LSS ID or you do not specify a value for the *dev*id variable in your profile file.

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 *dev* variable in your profile file.

The fully qualified LSS ID format is `storage_image_ID/XX`, where `XX` is a hexadecimal number in the range 00 - FE (2107 only).

If you specify the dash (-), this parameter information is automatically supplied.

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

Example of a shortened version: `00`

Example of multiple IDs: `10 20 30`

Example (2107)

Invoking the `unfreezeflash` command

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

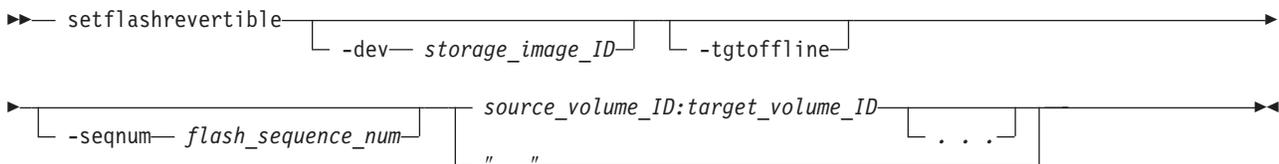
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

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

setflashrevertible

The `setflashrevertible` command 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, or reverted back to the last consistency. This command must be run before the FlashCopy pair can be committed or reverted.



Parameters

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

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the source volume and you do not specify a value for the *dev* variable in your profile file.

Example: `IBM.2107-75FA120`

-tgtoffline

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

-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 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* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the setflashrevertible command

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

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

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

Remote FlashCopy commands

This section contains commands that are used to configure Remote FlashCopy relationships and to display Remote FlashCopy information. 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.

Use the following commands to configure Remote FlashCopy relationships and to display Remote FlashCopy information. Also, see the *Processing Remote FlashCopy (inband) transactions* scenario for more details.

and serial number. This parameter is required if you do not specify a fully qualified ID for all the source volumes and you do not specify a value for the *dev* variable in your profile file.

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.

The fully qualified LSS ID format is *storage_image_ID/XX*, where *XX* is a hexadecimal number in the range 00 - FE (2107 only) .

Example: IBM.2107-75FA120/00

-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 System Storage 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 *dev* 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 *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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

Example of a shortened version: 0001

Example of IDs: 0001 0003 0008

Example (2107)

Invoking the commitremoteflash command

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

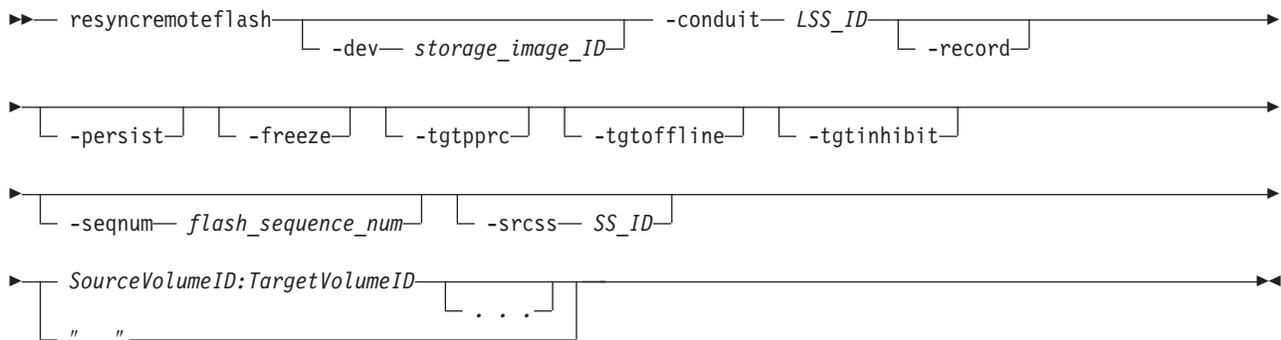
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

FlashCopy pair 0100:0200 successfully committed.

resyncremoteflash

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



Parameters

Note: When a pair is established with the **-record** and **-persist** parameters, the pair initially synchronizes and then a record of all data that is written from the host to the source is maintained in the source volumes. When the **resyncremoteflash** command is issued on the pair, the new data that is written to the source is copied to the target. The specified parameters in this command replace the parameters in the existing relationship. In order 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. This parameter is required if you do not specify a fully qualified ID for the source and target volumes or you do not specify a value for the *dev* variable in your profile file.

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.

The fully qualified LSS ID format is *storage_image_ID/XX*, where *XX* is a hexadecimal number in the range 00 - FE (2107 only).

Example: IBM.2107-75FA120/00

-record

(Optional) Specifies that the changed tracks on both volumes within a FlashCopy pair be recorded. Select this parameter if you intend to use the **resyncremoteflash** command again with this pair. The **-persist** parameter is automatically set in the code when the **-record** parameter is specified.

-persist

(Optional) Specifies that the FlashCopy relationship be retained after the background copy completes. The FlashCopy relationship between the source and the target volumes remains indefinitely until you issue a **rmremoteflash** command. This parameter is automatically set in the code when the **-record** parameter is specified.

Specify this parameter along with the **-record** parameter if you intend to use this volume pair with the **resyncremoteflash**, **reverseremoteflash**, or **setremoteflashrevertible** commands.

-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 in conjunction 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.

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: FF10

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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the **resyncremoteflash** command

```
dscli>resyncremoteflash -dev IBM.2107-75FA120  
-conduit IBM.2107-75FA150/10 0100:0200
```

The resulting output

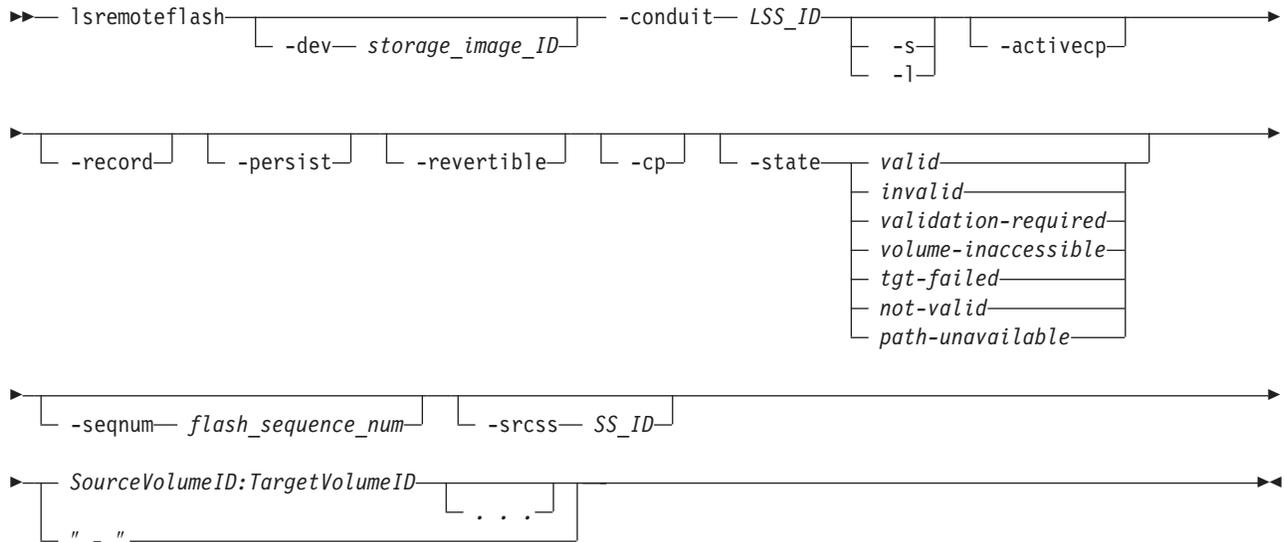
```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
Remote FlashCopy volume pair 0100:0200 successfully resynchronized.  
Use the lsremoteflash command to determine copy completion.
```

Note: This message is returned before the copy completes.

lsremoteflash

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



Parameters

Note: All settings apply to all FlashCopy pairs that are specified.

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which consists of manufacturer, machine type, and serial number. This parameter is required if you do not specify a fully qualified ID for the source and target volumes or you do not specify a value for the *dev* variable in your profile file.

-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 *XX* is a hexadecimal number in the range 00 - FE (2107 only).

-s (Optional) Specifies that FlashCopy pair IDs be displayed. You cannot use the **-l** and the **-s** parameters together.

-l (Optional) Specifies that the default detailed output plus Out-of-Sync Tracks and Date Created information be displayed. You cannot use the **-l** and the **-s** parameters together.

-activecp

(Optional) Specifies that FlashCopy relationships where their background copy process is active be displayed.

-record

(Optional) Specifies that the FlashCopy relationships that were established with the **-record** parameter be displayed.

-persist

(Optional) Specifies that the FlashCopy relationships that were established with the **-persist** parameter be displayed.

-revertible

(Optional) Specifies that the FlashCopy relationships that were established with the **-revertible** parameter be displayed.

-cp

(Optional) Specifies that the FlashCopy relationships that were established with the run background copy (**-cp**) parameter be displayed.

-state *valid | invalid | validation-required | volume-inaccessible | tgt-failed | not-valid | path-unavailable*

(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 of the *valid* state are displayed.

-seqnum *flash_sequence_number*

(Optional) Specifies that the FlashCopy relationships that are associated with the specified sequence number be displayed.

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 System Storage 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 source and target volume pairs with the specified IDs be displayed.

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.

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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 using the **-l** parameter. A separate example is not shown for the 1750 as the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

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

```
Date/Time: Sun Jan 09 02:23:49 PST 2005 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

ID	SrcLSS	Sequence Num	Active Copy	Recording	Persistent	Reversible
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
Enabled	Disabled	Disabled	Yes	0	12/01 /2003 02:20:00	12/01 /2003 02:23:47	Valid

Report field definitions

ID Specifies 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

Specifies the logical subsystem ID.

Sequence Num

Specifies the sequence number that is associated with the FlashCopy relationship.

ActiveCopy

Specifies (enabled or disabled) whether the background copy is active on the specified FlashCopy pair.

Recording

Specifies (enabled or disabled) whether the designated FlashCopy pair is established with recording activated.

Persistent

Specifies (enabled or disabled) whether the designated FlashCopy pair is established with persistent activated.

Revertible

Specifies (enabled or disabled) whether the designated FlashCopy pair is established with the revertible option activated.

SourceWriteEnabled

Specifies (enabled or disabled) whether or not this FlashCopy relationship was established with the Allow Source Writes option.

TargetWriteEnabled

Specifies (enabled or disabled) whether this FlashCopy relationship was established with the Allow Target Writes option

Background Copy

Specifies (enabled or disabled) whether this FlashCopy relationship was established with the Run Background Copy option.

OutofSync Tracks

Specifies the number of tracks that are not synchronized for this FlashCopy relationship.

DateCreated

Specifies the date and the time that the FlashCopy relationship was established.

DateSynced

Specifies the date and the time that this FlashCopy relationship was synchronized, or null (-) if the relationship is not synchronized.

State Specifies 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 null (-) 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. Generally, 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

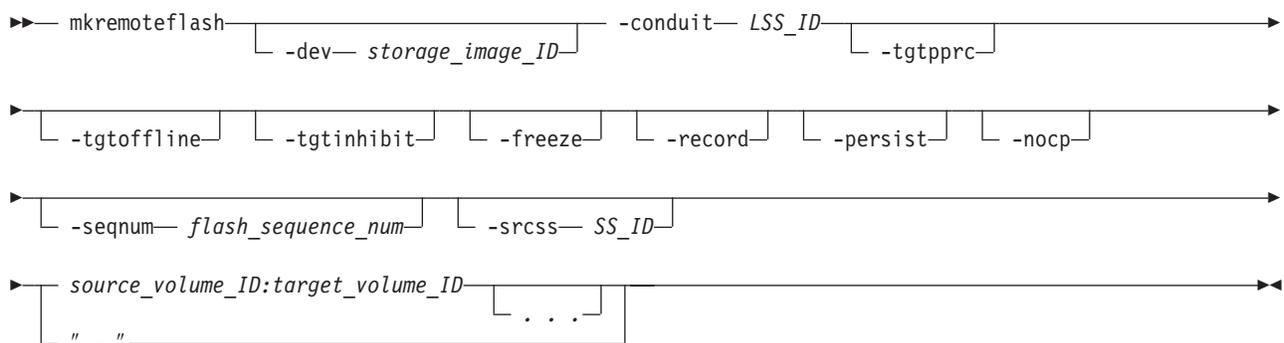
Indicates that a general internal error has occurred when the query is processed.

Path Unavailable

The specified inband path does not exist. The user should verify that the Remote Mirror and Copy pair exists.

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, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the source and target volumes or you do not specify a value for the *dev*id variable in your profile file.

-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.

The fully qualified LSS ID format is storage_image_ID/XX, where XX is a hexadecimal number in the range 00 - FE (2107 only).

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*. 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 in conjunction 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) Specifies that the changed tracks on both volumes within a FlashCopy pair be recorded. Select this parameter if you intend to use the **resyncremoteflash** command again with this pair. The **-persist** parameter is automatically selected when the **-record** parameter is specified.

-persist

(Optional) Specifies that you want to retain the FlashCopy relationship after the background copy completes. The FlashCopy relationship between the source and the target volumes remains indefinitely until you issue a **rmremoteflash** command. This parameter is automatically selected when the **-record** parameter is specified.

Specify this parameter along with the **-record** parameter if you intend to use this volume pair with the **resyncremoteflash**, **reverseremoteflash**, or **setremoteflashrevertible** commands.

-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.

-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 System Storage 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 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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the mkremoteflash command

```
dscli>mkremoteflash -dev IBM.2107-75FA120
-conduit IBM.2107-75FA150/10 0100:0200
```

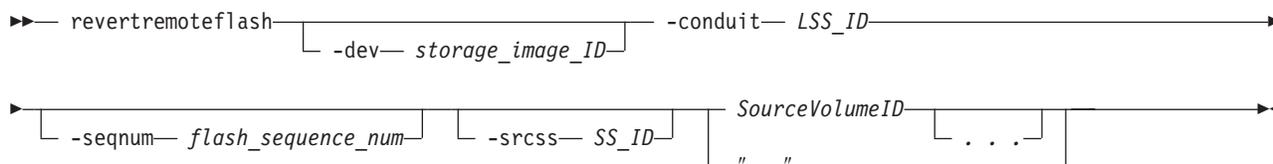
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

```
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, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the source volume or you do not specify a value for the *dev* variable in your profile file.

-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.

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 *XX* is a hexadecimal number in the range 00 - FE (2107 only).

Example: IBM.2107-75FA120/00

Example: IBM.1750-68FA120/00

-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 System Storage 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 *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

Example of a shortened version: 0001

Example of multiple IDs: 0001 0003 0008

Example (2107)

Invoking the revertremoteflash command

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

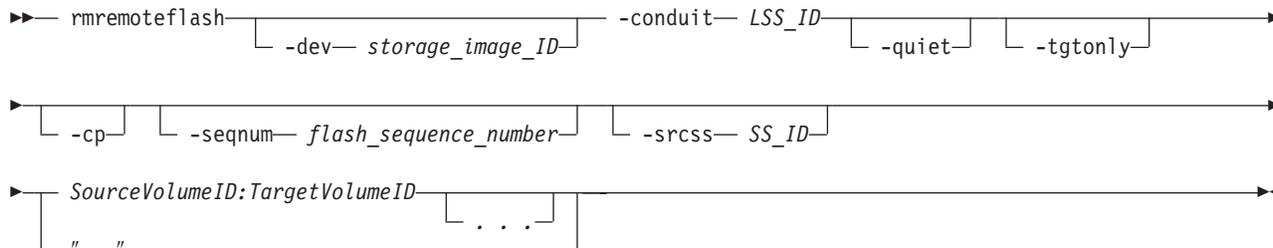
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

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

rmremoteflash

The **rmremoteflash** command removes a relationship between remote FlashCopy volume pairs.



Parameters

Notes:

1. Invoking this command and using the **-cp** parameter on a FlashCopy relationship that was previously marked with the **-persist** parameter does not remove the relationship. Instead, the source volume is copied to the target volume.
2. Invoking this command resets the **-tgtinhibit** parameter option if it was previously set.
3. All settings apply to all specified FlashCopy pairs.
4. The **-seqnum** parameter is not supported for model 2105.
5. When **SS_IDs** are specified, the source volumes are restricted to 1 LSS.

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the source and target volumes or you do not specify a value for the *devid* variable in your profile file.

-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*, where *XX* is a hexadecimal number in the range 00 - FE (2107 only).

-quiet

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

-tgtonly

(Optional) Specifies only the target volume of the remote FlashCopy pair that you want to remove. In addition, the *Copy Indicator* for the target volume is reset.

Note: You can use this parameter only with fixed block volumes.

-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 System Storage 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. When you use the **-tgtonly** parameter, enter volume IDs. Volume pair IDs cannot be used with the **-tgtonly** parameter.

The volume ID is a 32-bit number that can be represented as 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the **rmremoteflash** command

```
dscli>rmremoteflash -dev IBM.2107-75FA120  
-conduit IBM.2107-75FA150/10 0100:0200
```

The resulting output

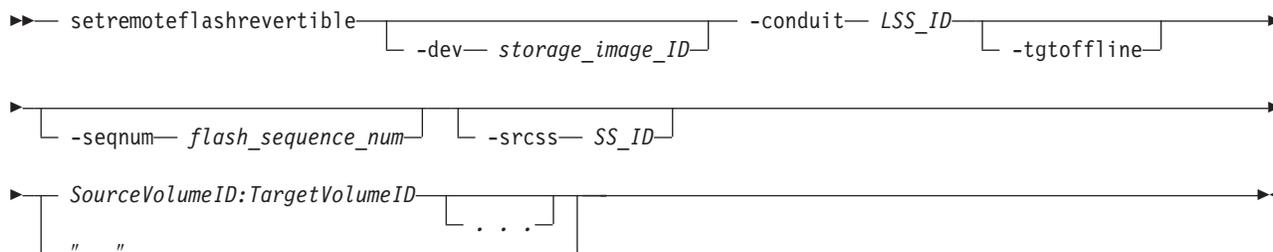
```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
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 Global Mirror relationship to *revertible*. This command must be run before the FlashCopy pair can be committed or reverted. Once a pair is revertible, the data can be committed to the target to form a new consistency, or reverted back to the last consistency.



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.

-dev *storage_image_ID*

(Optional) Specifies the storage image ID, which includes manufacturer, type,

and serial number. This parameter is required if you do not specify a fully qualified ID for the source volume or you do not specify a value for the *dev* variable in your profile file.

-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 passage 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 passage LSS volumes serves as a primary volume.

Accepts 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 (2107 only).

-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.

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: FF10

SourceVolumeID:TargetVolumeID . . . | -

(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 *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the `setremoteflashrevertible` command

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

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

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

Remote Mirror and Copy path commands

This section contains commands that create and delete Remote Mirror and Copy (formerly PPRC) paths and display Remote Mirror and Copy path information.

Use the following commands to create and delete Remote Mirror and Copy paths and to display Remote Mirror and Copy path information:

- `lsavailpprcport`
- `lspprcpath`
- `mkesconpprcpath`
- `mkpprcpath`
- `rmpprcpath`

The `lsavailpprcport` command 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.

The `lspprcpath` command generates a report that displays a list of existing Remote Mirror and Copy path definitions.

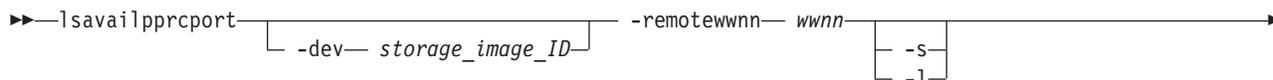
The `mkesconpprcpath` command creates a Remote Mirror and Copy path between source and target logical subsystems over an ESCON connection.

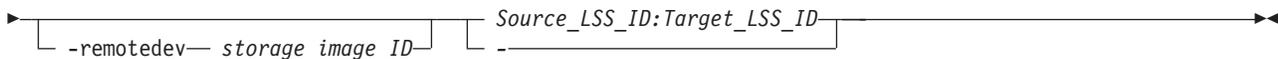
The `mkpprcpath` command establishes or replaces a Remote Mirror and Copy path between source and target logical subsystems (LSSs) over a fibre-channel connection.

The `rmpprcpath` command deletes one or more specified Remote Mirror and Copy paths.

`lsavailpprcport`

The `lsavailpprcport` command displays a list of ESCON or fibre-channel I/O ports that can be defined as remote mirror and copy (formerly PPRC) paths. The DS8000 supports only fibre-channel ports. The Enterprise Storage Server (2105 machine type) supports ESCON ports.





Parameters

-dev *storage_image_ID*

(Optional). Specifies the source volume storage image ID, which consists of manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the source logical subsystem.

Example: IBM.2107-75FA120

-remotewwnn *wwnn*

(Required). Specifies the worldwide node name (WWNN). The format is a 16-hexadecimal ID.

Note: You want to use the WWNN that is associated with the remote storage image. Run the **lssi** or **shows** command to obtain this number. If you use the WWNN that is associated with the primary storage unit, this command fails.

WWNN example: 12341234000A000F

-s (Optional). Displays the local port ID. You cannot use the **-l** and the **-s** parameters together.

-l (Optional). Displays all fields. You cannot use the **-l** and the **-s** parameters together.

-remotedev *storage_image_ID*

(Required or Optional). Specifies the remote storage unit that contains the I/O ports that are queried by the *Source_LSS_ID:Target_LSS_ID* parameter. The remotdev ID consists of the value for the manufacturer, machine type, and serial number.

Required - This parameter is required when querying ESCON I/O ports unless a fully qualified target logical subsystem ID is specified.

Optional - This parameter is optional if you are querying fibre channel I/O ports.

Note: If specified the format of this entry might be checked for correctness even though the value is not used.

Example: IBM.2107-75FA120

Source_LSS_ID:Target_LSS_ID | -

(Required). Queries I/O ports that are available for a remote mirror and copy path relationship for the source and target LSSs. This parameter accepts fully qualified LSS IDs, which includes the storage image ID or shortened version without the storage image ID, if the **-dev** parameter is specified.

A remote mirror and copy path LSS pair ID consists of two LSS IDs, one designated as the source LSS and the other as the target LSS for a remote mirror and copy path relationship. The two LSS IDs must be separated with a colon and no spaces. The first LSS ID is the source LSS. The second LSS ID is the target LSS.

The fully qualified LSS ID format is *storage_image_ID/xx*, where 'xx' is a hexadecimal number in the range 00 - FE (2107 only). For example, IBM.2107-75FA120/01:IBM.2107-75FA150/01

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. For example, your command line would look like the following:

```
dscli>lsavailpprcport -l -remotewwnn 12341234000A000F
IBM.2107-75FA120/01:IBM.2107-75FA150/01
```

If you specify the dash (-), this parameter information is automatically supplied.

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. A separate example is not shown for the 1750 because the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

Invoking the lsavailpprcport command

```
dscli>lsavailpprcport -l -dev IBM.2107-75FA120
-remotewwnn 12341234000A000F 01:01
```

The resulting output

ESCON port information displays for the 2105 machine type.

Local port	Attached port	Type	Switch ID	Switch port
I0100	I0200	FCP	N/A	N/A
I0150	I0620	ESCON	N/A	N/A
I0200	N/A	ESCON Switch	IBM.111.2222. 75113AB	I10
I0250	N/A	ESCON Switch	IBM.111.2222. 75113AB	I20

Report field descriptions

Local port

Specifies 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

Specifies 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 Specifies 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

Specifies the Switch ID for ESCON Switch connections.

Note: For FCP and direct ESCON, the displayed value in this field is **N/A** (not applicable).

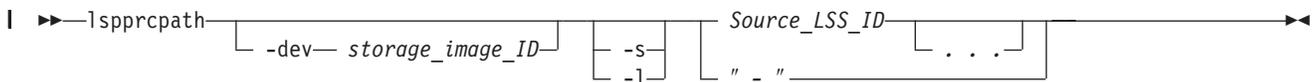
Switch port

Specifies the Port ID on the Switch device that is connected to the attached ESS. The Switch port ID component is two hexadecimal characters in the format 0xPP, where 'PP' is a port number (00 - ff). The number is prefixed with the letter I.

Note: For FCP and direct ESCON, the value of this field is **N/A** (not applicable).

lspprcpath

The **lspprcpath** command displays a list of existing Remote Mirror and Copy (formerly PPRC) 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. This parameter is required if you do not specify a fully qualified ID for the source LSS ID.

Example: IBM.2107-75FA120

-s (Optional) Displays the default output of the report but does not include the *Failed Reason* column. You cannot use the **-l** and the **-s** parameters together.

-l (Optional) Displays the default output and the *Failed Reason* descriptions. You cannot use the **-l** and the **-s** parameters together.

Source_LSS_ID . . . | -

(Required) Specifies that the Remote Mirror and Copy paths that are defined for the specified source LSS IDs be displayed.

This parameter accepts either a fully qualified LSS ID, 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 logical subsystem (LSS) hexadecimal number 00 - FE (for 2107).

You must separate multiple LSS IDs with spaces.

If you specify the dash (-), this parameter information is automatically supplied.

Example: 00

Example of multiple source LSS IDs: 00 01 02

Example

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

The following tables represent the headers that are displayed on the output report when the **-l** parameter is used with the **lspprcpath** command. A separate example is not shown for the 1750 because the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

Invoking the lspprcpath command

```
dscli>lspprcpath -dev IBM.2107-75FA120 -l 10
```

The resulting output

Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

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	Success	0012
IBM.2107 -75FA120 /10	IBM.2107 -75FA150 /13	Invalid	0013

Port	Attached Port	Tgt WWNN	Failed Reason
IBM.2107 -75FA120 /I0100	IBM.2107 -75FA150 /I0100	3007ACF 3012399E0	Configur ation Error
IBM.2107 -75FA120 /I0100	IBM.2107 -75FA150 /I0100	3007ACF- 3012399E0	-
IBM.2107 -75FA120 /I0100	IBM.2107 -75FA150 /I0100	3007ACF- 3012399E0	-

Port	Attached Port	Tgt WWNN	Failed Reason
IBM.2107 -75FA120 /I0100	IBM.2107 -75FA150 /I0100	3007ACF- 3012399E0	-

Report field definitions

- Src** Specifies the fully qualified logical subsystem ID.
- Tgt** Specifies the fully qualified Remote Mirror and Copy ID for the path target LSS.
- State** Displays the current Remote Mirror and Copy path state. One of the following values can be displayed:
- Failed** The transaction failed for one of the specified reasons. See the Failed Reason column for an explanation.
 - Success**
The transaction was a success. When this is the State value, the Failed Reason code displays a null (-) value.
 - Invalid**
The transaction was not valid. When this is the State value, the Failed Reason code displays a null (-).
- SS** Specifies the subsystem identifier (SSID) of the target LSS.
- Port** Specifies 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**
Specifies 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**
Specifies the worldwide node name of the remote storage image.
- Failed Reason**
Displays the reason for a failed transaction. One of the following values can be displayed:
- Configuration error**
A path has failed. Delete the original entry and resubmit the **mkpprcpath** command.
 - Down** An FCP path has failed because of a communication or hardware failure.
 - FCP source 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.

FCP target login exceeded

The maximum number of logins for each FCP path to the target LSS has been exceeded. The FCP target is unavailable.

FCP target unavailable

An FCP path to the target LSS is unavailable.

None The path has no state.

Source no resources

No resources are available at the source site for the logical paths to be established.

Retry Specifies the number of attempts to reestablish path connection.

Target mismatch

There is a target (remote) site system adapter ID mismatch.

Target no resources

Resources are not available at the target (remote) site to establish the logical paths.

Target LSS mismatch

The sequence number of the storage unit at the target site does not match the number of the target LSS.

Reserved

Not available for use at this time.

Timeout

A timeout has occurred. No reason is available.

Unknown

The state of the path is unknown.

Path Degraded Due to High Failure Rate

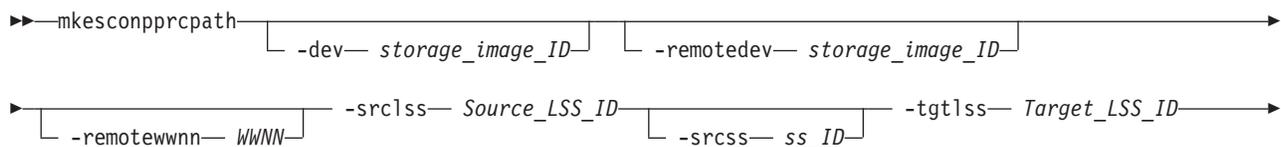
Indicates that a fibre channel path is established; however, because of the high failure rate, the path is degraded.

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.

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 System Storage Enterprise Storage Servers (2105, Model 800 and Model 750).





Parameters

Notes:

1. The **mkessconpprcpath** command is applicable **only** for the IBM System Storage Enterprise Storage Server (2105, Model 800 and Model 750). Models 2107 and 1750 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 source storage image ID, which consists of manufacturer, type, and serial number. This parameter is required if you do not fully qualify the source LSS ID.

Example: IBM.2105-75FA120

-remotedev *storage_image_ID*

(Optional). Specifies the remote storage image ID, which consists of manufacturer, type, and serial number. This parameter is required if you do not fully qualify the target LSS ID.

Example: IBM.2105-75FA150

-srcLSS *Source_LSS_ID*

(Required). Specifies the source logical subsystem (LSS) ID. Accepts a fully qualified LSS ID, which includes the storage image ID or a shortened version without the storage image ID, if the **-dev** parameter is used. The fully qualified LSS ID format is *storage_image_ID/xx*, where 'xx' is a hexadecimal number in the range '00 - FE'.

-srcSS *ss_ID*

(Optional). Specifies the subsystem ID of the primary logical subsystem in the format '0x0001 - 0xFFFF'.

This value is required for the IBM Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-tgtLSS *Target_LSS_ID*

(Required). Specifies the target logical subsystem (LSS) ID. Accepts a fully qualified LSS ID, which includes the storage image ID, or a shortened version without the storage image ID, if the **-remotedev** parameter is used. The fully qualified LSS ID format is *storage_image_ID/xx*, where 'xx' is a hexadecimal number in the range '00 - FE'.

-remotewwnn *WWNN*

(Optional). Specifies the worldwide node name. The format is a 16-hexadecimal ID.

Note: If you use this parameter, the format is checked even though there might be times that the value is not used.

Example: 12341234000A000F

-tgtss *ss_ID*

(Optional). Specifies the subsystem ID of the secondary logical subsystem in the format '0x0001 - 0xFFFF'.

This value is required for the IBM Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-consistgrp

(Optional). Creates a consistency group for the remote mirror and copy volume pairs that are associated with the PPRC paths that are established by this command. A remote mirror and copy consistency group is a set of remote mirror and copy volume pairs that have the same source and target LSS.

Normally, when an error occurs in a member of a remote mirror and copy volume pair, the volume is put in a *suspended* state. However, if the volume is participating in a consistency group, it is placed in a *long busy* state.

-force

(Optional). Creates a new remote mirror and copy path even if the specified remote mirror and copy path already exists.

Source_Port_ID:Target_Port_ID . . . | -

(Required). Establishes a remote mirror and copy path between the source and target ports for the specified source and target logical subsystems. The source port must be an ESCON I/O port that is configured for point-to-point or switch topology. The source port is enabled automatically for remote mirror and copy primary I/O operations. The target port must be a switch I/O port that is configured for point-to-point or switch topology. The target port is enabled automatically for remote mirror and copy primary I/O operations.

Note: Do not specify a target port ID when you specify an ESCON direct connection. Instead, specify only the source port ID.

This parameter accepts only nonfully 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.

If you specify the dash (-), this parameter information is automatically supplied.

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

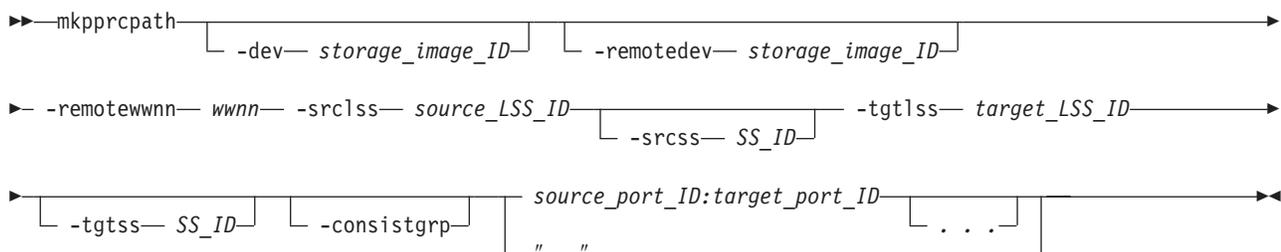
```
dscli>mkesconpprcpath -dev IBM.2105-75FA120 -remotedev IBM.2105-75FA150  
-srcLSS 01 -tgtLSS 01 I0100:I20 I0110:I21
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2105-75FA120  
Remote Mirror and Copy path  
IBM.2105-75FA120/01:IBM.2105-75FA150/01 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. 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.



Parameters

-dev *storage_image_ID*

(Optional) Specifies the source storage image ID, which includes manufacturer, machine type, and serial number. This parameter is required if you do not specify a fully qualified source port ID.

-remotedev *storage_image_ID*

(Optional) Specifies the ID of the secondary storage image, which includes manufacturer, machine type, and serial number. If specified, the format of this entry might be checked for correctness although the value is not used

Example: IBM.2107-75FA150

-remotewwnn *wwnn*

(Required) Specifies the worldwide node name of the secondary storage image. The format is a 16-hexadecimal ID.

Note: Ensure that you use the worldwide node name that is associated with the secondary storage unit. Run the **lssi** or **showsi** command to obtain this number.

Example: 12341234000A000F

-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 (2107) 00 - FE.

Example of a shortened version: 00

-srcss *SS_ID*

(Optional) Specifies the subsystem ID of the primary logical subsystem in the format 0x0001 - 0xFFFF.

This value is required for the IBM TotalStorage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-tgtlss *target_LSS_ID*

(Required) Specifies the logical subsystem ID associated with the secondary storage unit 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 (2107) 00 - FE.

Example of a shortened version: 01

-tgtss *SS_ID*

(Optional) Specifies the subsystem ID of the secondary logical subsystem in the format 0x0001 - 0xFFFF.

This value is required for the IBM TotalStorage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-consistgrp

(Optional) Creates a consistency group for the remote mirror and copy volume pairs. A remote mirror and copy consistency group is a set of remote mirror and copy volume pairs that have the same source and target LSS.

Normally, when an error occurs in a member of a remote mirror and copy volume pair, the storage unit 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 includes 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.

This parameter accepts up to eight remote mirror and copy path port pair IDs. You must separate multiple port pair IDs with spaces.

If you specify a dash (-), this parameter information is automatically supplied.

Example of the shortened version: I1A10:I2A20

Example of multiple pairs: I1A10:I2A20 I1A11:I2A21 I1A12:I2A22

Example (2107)

Invoking the mkpprcpath command

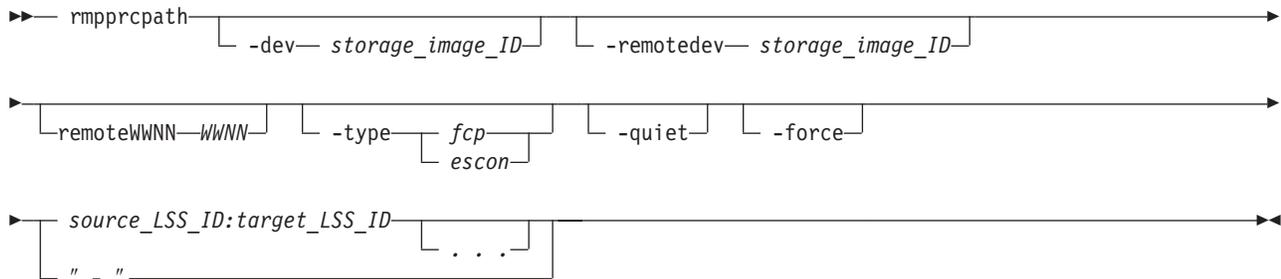
```
dscli>mkpprcpath -dev IBM.2107-75FA120  
-srcLSS 01 -tgtLSS 01 -remotewwn 12341234000A000F I0100:I0100
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120  
Remote Mirror and Copy path 01:01 successfully established.
```

rmpprcpath

The rmpprcpath 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. This parameter is required if you do not specify a fully qualified source LSS ID or you do not specify a value for the *devId* variable in your profile file.

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.

-remoteWWNN *WWNN*

(Optional) Specifies the secondary worldwide node name.

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 *fcpl escon*

(Optional) The type of the connection over which the path was created.

fcpl Fibre-channel protocol

escon Enterprise Systems Connection (IBM S/390 and zSeries)

-quiet

(Optional) Turns off the 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 (2107) 00 - FE.

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the `rmpprcpath` command

```
dscli>rmpprcpath -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150
-remotewwnn 12341234000A000F 01:01
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

```
Are you sure want to remove the Remote Mirror and Copy path 01:01? [y/n]: Y
Remote Mirror and Copy path 01:01 successfully removed.
```

Remote Mirror and Copy pair commands

This section contains commands that create, manage, and delete Remote Mirror and Copy (formerly PPRC) pairs and display Remote Mirror and Copy pair information.

Use the following commands to create, manage, and delete Remote Mirror and Copy pairs and to display Remote Mirror and Copy pair information:

- **failbackpprc**

- **failoverpprc**
- **lspprc**
- **mkpprc**
- **freezepprc**
- **pausepprc**
- **resumepprc**
- **rmpprc**
- **unfreezepprc**

The **failbackpprc** command copies the required data from the source volume to the target volume in order 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).

The **failoverpprc** command 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.

The **lspprc** command 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.

The **mkpprc** command establishes a remote mirror and copy relationship for a volume pair.

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/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.

The **pausepprc** command pauses an existing remote mirror and copy volume pair relationship or pauses a single volume ID.

The **resumepprc** command resumes a remote mirror and copy relationship for a volume pair.

The **rmpprc** command 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).

-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 that this parameter reflect a role reversal. The value for this parameter must be the **original secondary** site. For example:

- Original primary site (Site A) has a value of (2107) IBM.2107-75FA120 with volumes 0100, 0101, 0102, 0103.
- Original secondary site (Site B) has a value of (2107) IBM.2107-75FA150 with volumes 0200, 0201, 0202, 0203.
- The following **failbackpprc** command is correct:

```
(2107) 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.

-suspend

(Optional) Specifies that the remote mirror and copy relationship 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 be established, including when the target volume is online to host systems.

Note: This parameter applies only to S/390 or zSeries volumes. It does not apply to Open Systems volumes.

-resetreserve

(Optional) Specifies that a remote mirror and copy relationship 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 those described in 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.

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Note: The source and target volumes are restricted to one LSS for the source and one LSS for the target when you specify *ss_IDs*.

Example: 0010

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Note: The source and target volumes are restricted to one LSS for the source and one LSS for the target when you specify *ss_IDs*.

Example: 0010

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* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

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

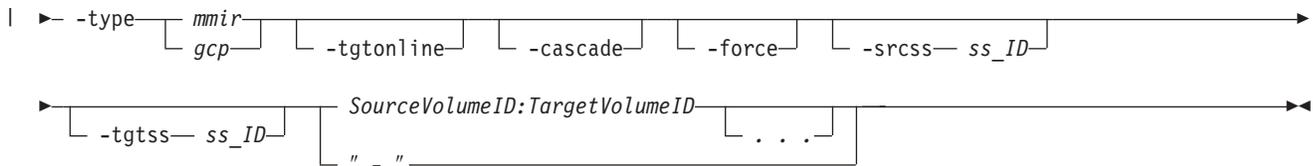
```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107.75FA120
```

```
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 using 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_ID ] [ -remotedev— storage_image_ID ]
```



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 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 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, type, and serial number. This parameter is required if you do not specify a fully qualified source volume ID.

Note: The use of the **failoverpprc** command requires that this parameter reflect a role reversal. The value for this parameter must be the **original secondary** site. For example:

- Original primary site (Site A) has a value of (2107) IBM.2107-75FA120 with volumes 0100, 0101, 0102, 0103.
- Original secondary site (Site B) has a value of (2107) IBM.2107-75FA150 with volumes 0200, 0201, 0202, 0203.
- The following **failoverpprc** command is correct:

```

(2107) 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 that this parameter reflect a role reversal. The value for this parameter must be the **original primary** site. For example:

- Original primary site (Site A) has a value of (2107) IBM.2107-75FA120 with volumes 0100, 0101, 0102, 0103.
- Original secondary site (Site B) has a value of (2107) IBM.2107-75FA150 with volumes 0200, 0201, 0202,0203.
- The following **failoverpprc** command is correct:

```
(2107) 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 performed at a later time. 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 S/390 or zSeries 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.

-force

(Optional) Specifies whether validation of the volumes that are involved in the establish request occurs or is bypassed. This parameter allows you to create a FlashCopy pair between two volumes that 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. 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 in this guide that describe a set of circumstances that allow you to safely use this parameter. If your circumstances do not match those described in the scenarios, you are cautioned not to use this parameter unless advised to do so by IBM Technical Support.

|

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

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 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 *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE, and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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 (2107)

Invoking the failoverpprc command

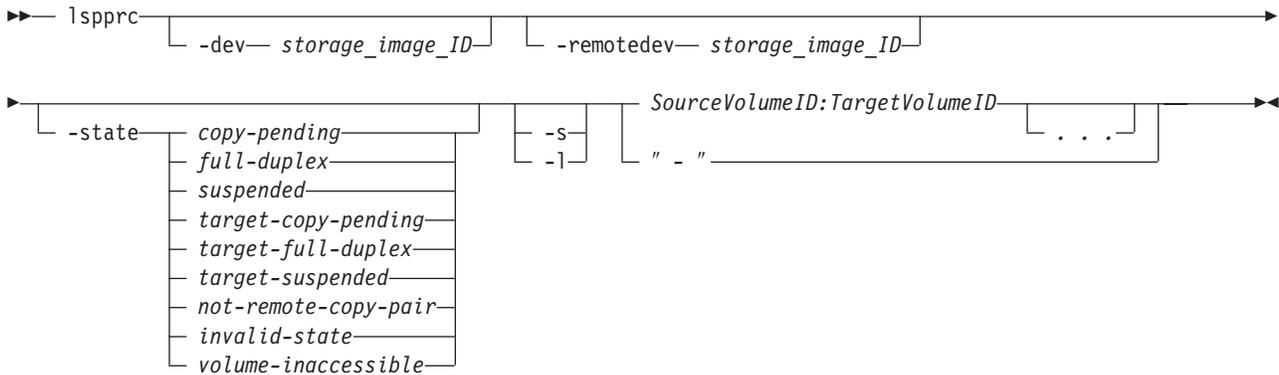
```
dscli>failoverpprc -dev IBM.2107-75FA150 -remotedev IBM.2107-75FA120  
0200:0100 0201:0101 0202:0102
```

The resulting output

A confirmation message is presented for each remote mirror and copy pair that is successfully suspended.

lspprc

The **lspprc** command displays a list of remote mirror and copy (formerly PPRC) volume relationships for a storage image, and 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. This parameter is required if you do not specify a fully qualified ID for the source and target volumes.

Example: IBM.2107-75FA120

-remotedev *storage_image_ID*

(Optional most times, 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.

Example: IBM.2107-75FA150

-state *copy-pending* | *full-duplex* | *suspended* | *target-copy-pending* | *target-full-duplex* | *target-suspended* | *not-remote-copy-pair* | *invalid-state* | *volume-inaccessible*

(Optional). Identifies the state of the remote mirror and copy relationship 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.

not-remote-copy-pair

Specifies that you want to view remote mirror and copy relationships that are not a remote copy pair.

invalid-state

Specifies that you want to view remote mirror and copy relationships that do not meet requirements for being valid. The report that is generated with this query only displays 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 null (-) value.

volume-inaccessible

Specifies that you want to view remote mirror and copy relationships where the volume cannot be viewed, generally this means that the volume is fenced. The report that is generated with this query only displays 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 null (-) value.

- s (Optional). Displays the remote mirror and copy volume pair IDs. You cannot use the -I and the -s parameters together.
- I (Optional). Displays the default output and the *target read enabled*, *source cascade*, *target cascade*, and *suspended* attributes. You cannot use the -I and the -s parameters together.

SourceVolumeID:TargetVolumeID . . . | -

(Required) Displays the remote mirror and copy relationships for the source and target volume pairs with the specified 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. 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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

Example of a shortened version: 0100:0100

Example of multiple IDs: 0100:0100 0200:0200 0300:0300

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 using the **-l** parameter. A separate example is not shown for the 1750 because the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

Invoking the lspprc command

```
dscli>lspprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 -l 0100:0100
0101:0101
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

ID	State	Reason	Type	Out Of Sync Tracks	Tgt Read	Src Cascade
IBM.2107-75FA120 /0100: IBM.2107-75FA150 /0100	Full-Duplex	-	Metro Mirror	0	Enabled	Disabled
IBM.2107-75FA120 /0101: IBM.2107-75FA150 /0101	Full-Duplex	-	Metro Mirror	0	Enabled	Disabled

Tgt Cascade	Date Sus pended	Source LSS	Timeout (secs)	Crit Mode	First Pass Status	GMIR CG	PPRC CG
Enabled	-	01	120	Disabled	False	Disabled	Enabled
Enabled	-	02	120	Disabled	False	Enabled	Disabled

Tgt Cascade	Date Suspended	Source LSS	Timeout (secs)	Crit Mode	First Pass Status	Incremental Resync	Tgt Write
Enabled	-	01	120	Disabled	False	Enabled	Enabled
Enabled	-	02	120	Disabled	False	Enabled	Enabled

Report field descriptions

ID Specifies the source and target volume IDs of a remote mirror and copy volume relationship.

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.

Not Part of Remote Copy Pair

Indicates that the relationship is not a remote copy pair.

Invalid State

Indicates that the listed remote mirror and copy volume relationships do not meet the requirements for being valid.

Note: The report that is generated with the *invalid-state* designation only displays 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 null (-) value.

Volume Inaccessible

Indicates that the volume could not be queried. Generally, this indicates that the volume is fenced.

Note: The report that is generated with the *Volume Inaccessible* designation only displays 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 null (-) value.

Reason

Indicates why the Remote Mirror and Copy volume relationship is suspended. The following values can be displayed:

Unknown

Specifies that the volume is suspended but the reason for the suspension is not specifically defined within the system

Not in PPRC Relationship

Specifies that the designated volume is not part of a Remote Mirror and Copy pair.

Host Source

Specifies that the Remote Mirror and Copy processing on the volume was suspended by the primary host. The host command might have specified an immediate suspension or that the volume be suspended when it entered a full duplex state.

Host Target

Specifies 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

Specifies 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

Specifies 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

Specifies that Remote Mirror and Copy processing was suspended on a volume when the secondary control unit sent a state change interrupt to the primary control unit indicating a transition to a simplex state.

Internal Conditions Target

Specifies 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 Specifies 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. Issue the **lspprc** command and use the generated report to determine the path status.

- b. Issue the **mkpprc** command if the paths are still in tact. This process resynchronizes the volume pairs.
 - c. Continue with your processing.
2. If the above process cannot be performed, 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 Specifies that Remote Mirror and Copy processing was suspended on a volume pair because the host issued a Freeze PPRC Group order.

Volume Not Configured

Specifies that Remote Mirror and Copy processing was suspended on a volume because the volume is not part of a copy pair relationship.

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 is dependent on the source volume size.

Notes:

1. If you issue 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 issue 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 issue the **lspprc** command and reissue it again before the 10 seconds has expired, there is no observable change in the value.
3. If you issue the **lspprc** command to view the out of sync value for a volume pair or a single volume on a 2107, there is an observable decrease in the value but only at 10 second intervals.

Tgt Read

Indicates that Read IO operations to the target volume are allowed.

Src Cascade

Indicates that the source volume of this relationship is enabled to also be a target volume of a different relationship.

Tgt Cascade

Indicates that the target volume of this relationship is enabled so that it is also a source volume for a different relationship.

Date Suspended

Indicates the date when this relationship was last suspended. The value can be displayed as a null (-).

SourceLSS

Indicates the consistency group LSS ID that is associated with the source volume of this PPRC volume relationship.

Timeout (secs)

Indicates the consistency group Long Busy Timeout setting for the LSS ID that is associated with the source volume of this PPRC volume relationship. The timeout value is either the default value of 120 seconds, or a user-specified value of 1 to 600 000 seconds.

Crit Mode

Indicates whether the remote copy and mirror primary volume represents a critical volume.

First Pass Status

Indicates the first pass Global Copy state of the source volume. The data is valid only when you query the primary box and the queried pair is a Global Copy pair.

Incremental Resync

Indicates if incremental resynchronization is running.

Tgt Write

Indicates whether input is allowed to the remote mirror and copy secondary volume.

GMIR CG

Indicates if 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 **-I** parameter and when the primary volume is being queried.
2. This value is not reported for model 2105. If a model 2105 is being queried, the value displayed for this field is N/A (not available).

PPRC CG

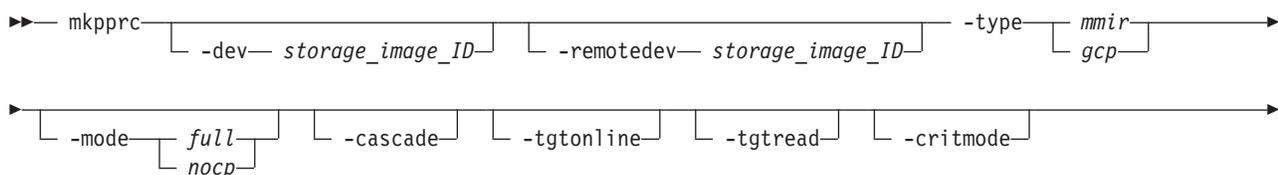
Indicates if the volume is in consistent asynchronous remote mirror and copy mode or is not available. The values that can be displayed are enabled, disabled, or not available.

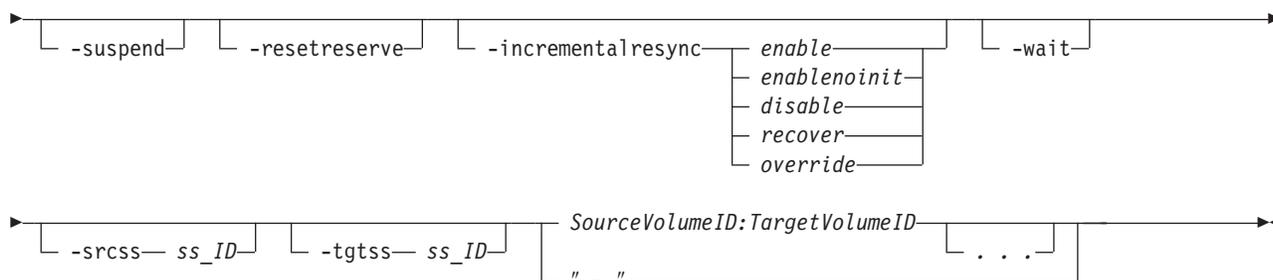
Notes:

1. This value is displayed when you designate the use of the **-I** parameter.
2. This value is not reported for model 2105. If a model 2105 is using being queried, the value displayed for this field is N/A (not available).

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. This parameter is required if you do not specify a fully qualified source volume ID.

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.

Example: IBM.2107-75FA120

-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 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 is the default when you do not specify no copy.

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 S/390 or zSeries 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 IBM S/390 or zSeries 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 S/390. The default value for this parameter is disabled.

Note: This parameter applies only to S/390 or zSeries volumes.

Critical mode operates in one of three ways depending on the setting of the LCU critical mode and the setting of the **-critmode** parameter in this command. 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	<ul style="list-style-type: none"> Suspends the primary volume. Allows write operations to the primary volume.
Critical Volume	Disabled	Enabled	<ul style="list-style-type: none"> Suspends the primary volume when the last path to the secondary volume has failed. Inhibits write operations to the primary volume.
Critical Heavy	Enabled	Enabled	<ul style="list-style-type: none"> Suspends the primary volume when the secondary volume cannot be updated for any reason. Inhibits write operations to the primary volume.

Notes:

1. Use the **-critmode** parameter only for log devices, not for devices that the system requires. In extreme cases, the host system might have to be IPLed in order 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 issue 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. This parameter can only be used 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 used in a three-site Metro/Global Mirror disaster recovery configuration. A three-site Metro/Global Mirror configuration usually involves: 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, and Site C (remote site), which contains the C volumes and an asynchronous copy located remotely from sites A and B.

You can specify the following options when you first establish volume pairs using the **mkpprc** command or on already established volume pairs using 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 flight" in a Metro/Global Mirror configuration. (In flight data refers to data that has been written to the primary volumes but has not been 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. This option should only be used 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 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 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 processing the **mkpprc** command. 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.

-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 issue a new command.
3. If you are running in single-shot mode, you can periodically issue the **lspprc** command to check the remote mirror and copy volume pair state, and then continue with new commands when the correct state is reached.

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

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 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 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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

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

Additional 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 circumstance might be processed, and both are correct. The first method might be fine if you were managing a few volumes, while the second allows you to process 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 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 (2107)

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)

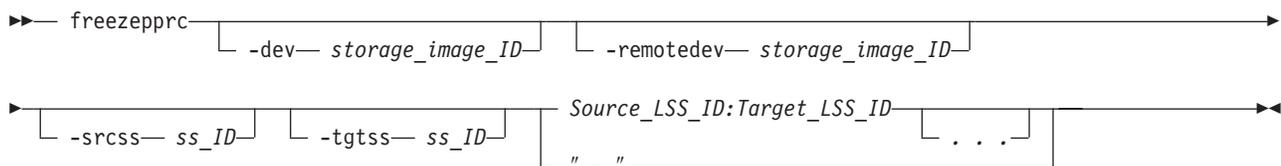
```
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
2100:2100 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
2101:2101 successfully created.
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
2102:2102 successfully created.
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.
```

freezepprc

The **freezepprc** command creates a new remote mirror and copy consistency group. It places the source logical subsystem (LSS) in the *long busy* state so that no I/Os can be directed to it. It also removes remote mirror and copy paths between the source LSS and target LSS and sets the *queue full* condition for the primary volume. This causes the host to queue writes to the primary volume until the *queue full* condition is reset. During the *queue full* condition, the primary volume reports *long busy* status.



Parameters

Note: When specifying SSIDs, the command is limited to one LSS pair.

-dev *storage_image_ID*

(Optional). Specifies the source storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified source LSS ID.

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.

Example: IBM.2107-75FA150

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

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 (2107 only) .

If you specify the dash (-), this parameter information is automatically supplied.

Example of a pair: 00:00

Example of multiple pairs: 00:00 01:01 02:02

Example (2107)

Invoking the freezepprc command

```
dscli>freezepprc -dev IBM.2107-75FA120  
-remotedev IBM.2107-75FA150 01:01
```

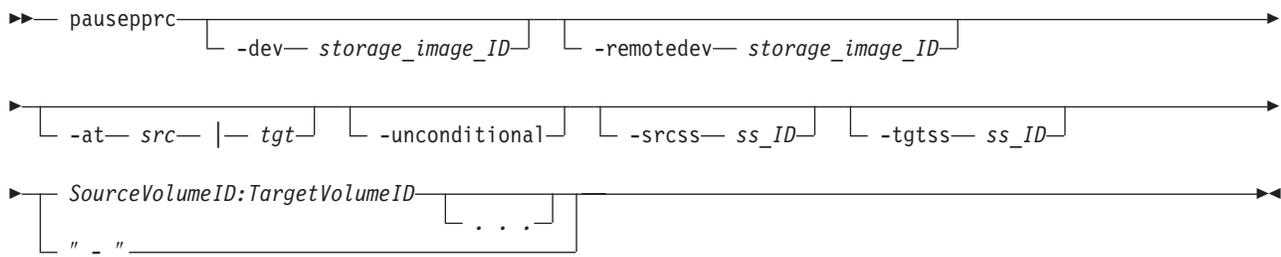
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
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, this command can be used to pause a single volume ID. To use 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 specifying SSIDs, the command is limited to one LSS pair.

-dev *storage_image_ID*

(Optional). Specifies the source storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified source volume ID or you do not specify a value for the *devid* variable in your profile file.

Example: IBM.2107-75FA120

-remotedev *storage_image_ID*

(Optional most times, however required as noted). Specifies the target storage image ID, which includes manufacturer, type, and serial number.

Note: The **-remotedev** parameter is required when volume pairs are specified and the **-dev** parameter is specified as well.

Example: IBM.2107-75FA120

-at *src | tgt*

(Optional).

src Select the **-at source** parameter option to initiate a suspend action from the source volume. After the task successfully runs, the source and target volumes are in the *suspend* state.

The **-at source** 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 Select the **-at target** parameter option to initiate 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 target** 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.

-unconditional

(Optional). Use this parameter to indicate 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. When you specify the **-unconditional** parameter, you can specify only the source volume ID if the **-at source** parameter option is selected, or specify only the target volume ID if the **-at target** parameter option is selected. Do not use volume pair IDs.

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

SourceVolumeID:TargetVolumeID . . . | -

(Required). Specifies that a remote mirror and copy volume relationship for the source and target volume pairs with the IDs specified 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 4 hexadecimal digits in the form of *XYZZ* where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

Note: Requests directed to target volumes are sent to the Storage Image identified by the **-dev** parameter or embedded in the fully qualified single volume IDs.

Example of a fully qualified pair: IBM.2107-75FA120/0100:IBM.2107-75FA150/0100

Example of a shortened version: 0100:0100

Example of multiple pairs: 0101:0101 0102:0102 0103:0103

Example (2107)

Invoking the pausepprc command

```
dscli>pausepprc -dev IBM.2107-75FA120  
-remotedev IBM.2107-75FA150 0100:0100
```

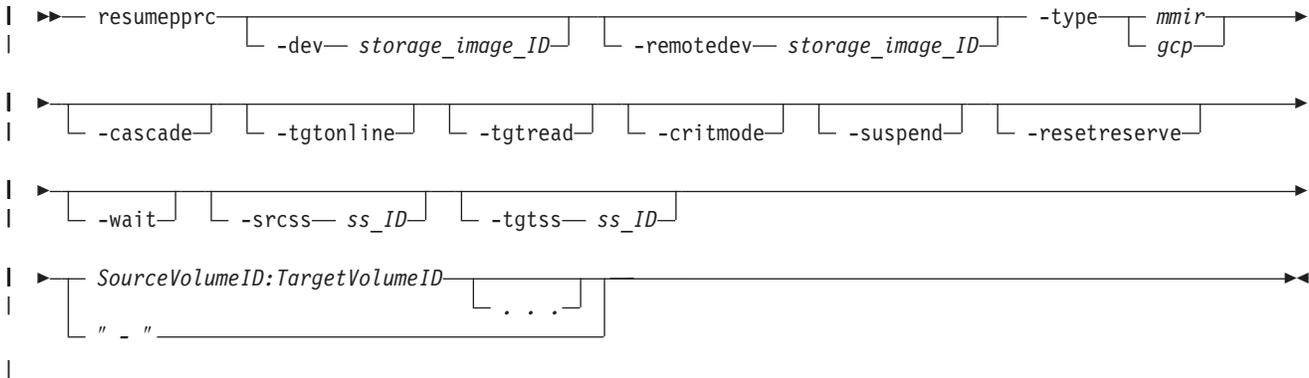
The resulting output

Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

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 1 LSS for the source and 1 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. This parameter is required if you do not specify a fully qualified source volume ID.

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 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 S/390 or zSeries volumes. It does not apply to Open Systems volumes.

-tgtread

(Optional) Allows host servers to read from the remote mirror and copy target volume. For a host server to read the volume, the remote mirror and copy pair must be in a full-duplex state. This parameter applies to open systems volumes and does not apply to IBM S/390 or zSeries 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 S/390.

-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.

-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.

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

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 *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE, and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the resumepprc command

```
dscli>resumepprc -dev IBM.2107-75FA120  
-remotedev IBM.2107-75FA150 0100:0100
```

The resulting output

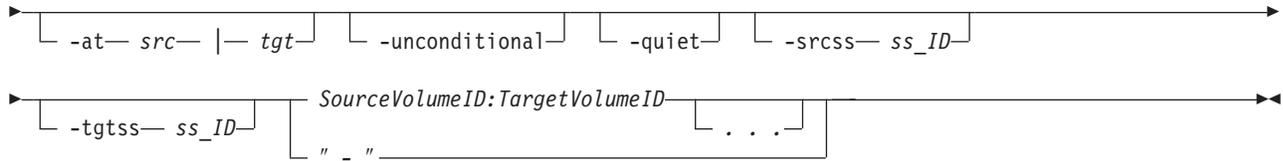
```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
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, this command can be used to remove 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). To use 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. The **-unconditional** parameter must be specified when you designate a single volume; otherwise an error occurs and the command process fails.

```
►►— rmpprc —┬── -dev— storage_image_ID —┬── -remotedev— storage_image_ID —┬──
```



Parameters

Notes:

1. When you specify subsystem IDs, the source and target volumes are restricted to 1 LSS for the source and 1 LSS for the target.
2. If there is a communication problem between the primary server and the secondary server (2-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 this circumstance has occurred, reissue the **rmpprc** command for each volume that did not have its pair relationship removed. This transaction involves using the **-at src** or **-at tgt** parameter and the **-unconditional** parameter.

The following 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 SourceVolumeID
```

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 3-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. This parameter is required if you do not specify a fully qualified source volume ID or you do not specify a value for the *devid* variable in your profile file.

Example: IBM.2107-75FA120

-remotedev *storage_image_ID*

(Optional most times, however required as noted). 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.

Example: IBM.2107-75FA150

-at *src | tgt*
(Optional).

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.

-unconditional

(Optional). Use this parameter to indicate 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. When you specify this parameter, you can specify only the source volume ID if the **-at src** parameter is selected, or specify only the target volume ID if the **-at tgt** parameter is selected. Do not use volume pair IDs.

-quiet

(Optional). Specifies that the confirmation prompt for this command be turned off.

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

SourceVolumeID:TargetVolumeID . . . | -

(Required). Specifies the Remote Mirror and Copy volume relationship for the source and target volume pairs with the IDs specified that is to be removed.

Note: Provide a single volume ID instead of a volume pair if you use the **-unconditional** parameter. Specifying pairs will result 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 4 hexadecimal digits in the form of *XYZZ*, where *X* is the address group (0 - 1), *XY* together is the logical subsystem number 00 - FE (for 2107), and *ZZ* is the volume number (00 - FF).

If you specify the dash (-), this parameter information is automatically supplied.

Example (2107)

Invoking the `rmpprc` command

```
dscli>rmpprc -dev IBM.2107-75FA120 -remotedev IBM.2107-75FA150 0100:0100
```

The resulting output

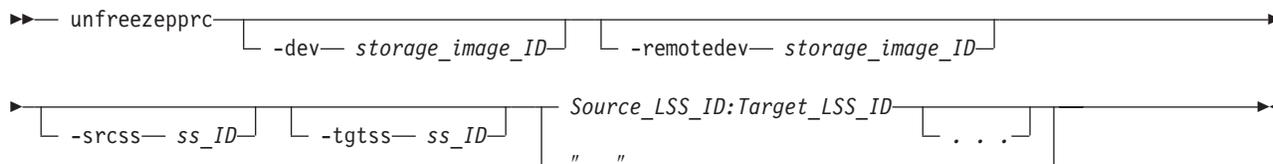
```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

```
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.



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. This parameter is required if you do not specify a fully qualified source LSS ID or you do not specify a value for the *dev* variable in your profile file.

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.

Example: IBM.2107-75FA120

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

-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 System Storage Enterprise Storage Server versions 2.4.0 and 2.4.1.

Example: 0010

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 (2107 only).

If you specify the dash (-), this parameter information is automatically supplied.

Example pair: 00:00

Example of multiple pairs: 00:00 01:01 02:02

Example (2107)

Invoking the `unfreezepprc` command

```
dscli>unfreezepprc -dev IBM.2107-75FA120  
-remotedev IBM.2107-75FA150 01:01
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
Remote Mirror and Copy pair ID 01:01  
successfully thawed.
```

Global Mirror commands

This section contains commands that are used to create, manage, and delete Global Mirror relationships and to display Global Mirror information.

Use the following commands to create, manage, and delete Global Mirror relationships and to display Global Mirror information:

- **mkgmir**
- **pausegmir**
- **resumegmir**
- **rmgmir**
- **showgmir**
- **showgmircg**
- **showgmiroos**

The **mkgmir** command starts Global Mirror processing for a specified session.

The **pausegmir** command 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

The **resumegmir** command 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.

The **rmgmir** command ends Global Mirror processing for the specified session.

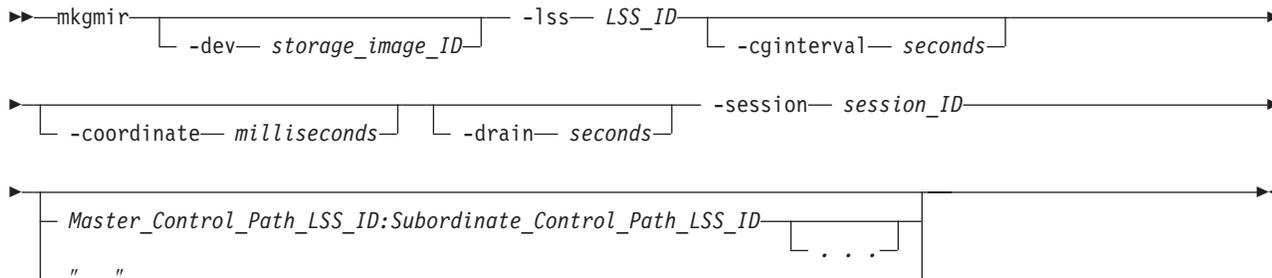
The **showgmir** command generates two reports. The first report displays the detailed properties about the current Global Mirror operations that are associated with a specified logical subsystem ID. The second report displays the performance metrics for the current Global Mirror operations that are associated with a specified logical subsystem ID.

The **shogmircg** command generates a report that displays the consistency group status for the specified Global Mirror session.

The **showgmiroos** command 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, type, and serial number. This parameter is required if you do not specify a fully qualified LSS ID, which includes a storage image ID.

Example: IBM.2107-75FA120

-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.

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

-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 or the 1 - 255 decimal 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 two characters in the 00 - FE (2107 only) hexadecimal range. 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.

If you specify the dash (-), this parameter information is automatically supplied.

Example of one Global Mirror association with a single subordinate in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00

Example of multiple Global Mirror associations with two subordinates in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00
IBM.2107-75FA120/00: IBM.2107-75FA150/01

Example (2107)

Invoking the mkgmir command

```
dscli>mkgmir -dev IBM.2107-75FA120 -lss 10  
-session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00
```

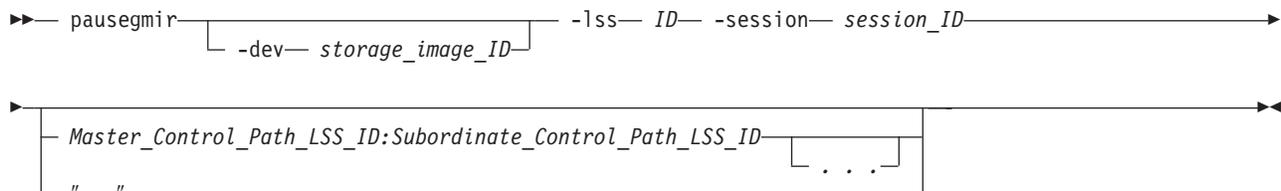
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

Global Mirror for session 01 successfully started.

pausegmir

The **pausegmir** command pauses Global Mirror for the specified session.



Parameters

-dev *storage_image_ID*

(Optional). Specifies the storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified LSS ID.

Example: IBM.2107-75FA120

-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.

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

-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 or the 1 - 255 decimal 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 (2107 only) . 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.

If you specify the dash (-), this parameter information is automatically supplied.

Example of one Global Mirror association with a single subordinate in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00

Example of multiple Global Mirror associations with two subordinates in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00
IBM.2107-75FA120/00: IBM.2107-75FA150/01

Example (2107)

Invoking the **pausegmir** command

```
dscli>pausegmir -dev IBM.2107-75FA120 -lss 10  
-session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00
```

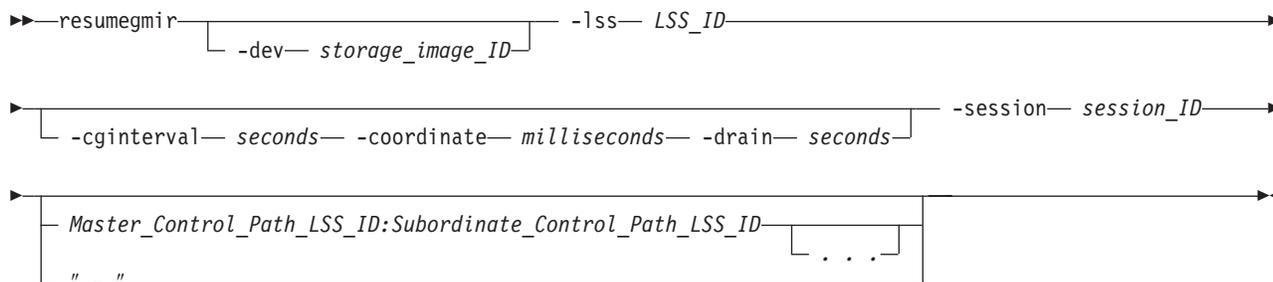
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

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. This parameter is required if you do not specify a fully qualified LSS ID or you do not specify a value for the *devid* variable in your profile file.

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 (2107 only) .

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. Your command would look like the following:

```
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 or the 1 - 255 decimal 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 (2107 only) . 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.

If you specify the dash (-), this parameter information is automatically supplied.

Example of one Global Mirror association with a single subordinate in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00

Example of multiple Global Mirror associations with two subordinates in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00
IBM.2107-75FA120/00: IBM.2107-75FA150/01

Example (2107)

Invoking the resumegmir command

```
dscli>resumegmir -dev IBM.2107-75FA120 -lss 10  
-session 01 IBM.2107-75FA120/00:IBM.2107-75FA150/00
```

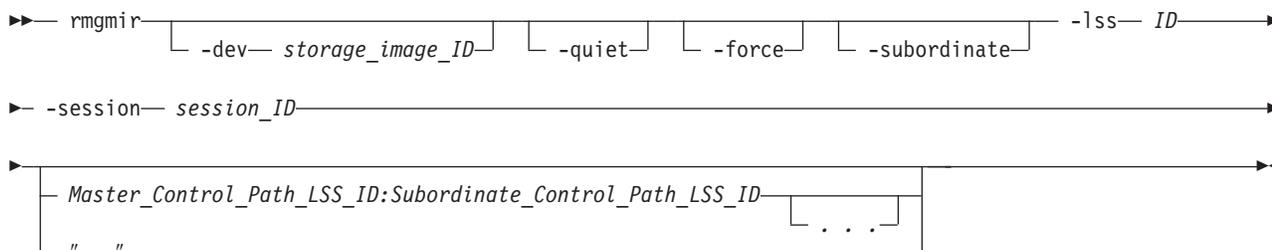
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

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, type, and serial number. This parameter is required if you do not specify a fully qualified LSS ID.

Example: IBM.2107-75FA120

-quiet

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

-force

(Optional). Forces the Global Mirror process to stop regardless of the state of the Global Mirror associations.

-subordinate

(Optional). Indicates that the **-lss** parameter specifies a subordinate LSS ID.

-lss *ID*

(Required). Specifies the logical subsystem (LSS) that is participating in the Global Mirror session. Accepts a fully qualified LSS ID, which includes the storage image ID or a shortened version without the storage image ID, if the **-dev** parameter is specified.

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

-session *session_ID*

(Required). Specifies the session ID for which the Global Mirror path association will be removed. A session ID is a Global Mirror session number that you assign in the 01 - FF hexadecimal or 1 - 255 decimal ranges.

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 (2107 only) . 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.

If you specify the dash (-), this parameter information is automatically supplied.

Example of one Global Mirror association with a single subordinate in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00

Example of multiple Global Mirror path associations with two subordinates in the configuration: IBM.2107-75FA120/00: IBM.2107-75FA150/00
IBM.2107-75FA120/00: IBM.2107-75FA150/01

Example (2107)

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

Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

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 detailed properties and performance metrics for a Global Mirror logical subsystem ID.

```
▶— showgmir ————— LSS_ID —▶  
    └─ -dev— storage_image_ID ┘ └─ -metrics ┘ └─ " - " ┘
```

Parameters

-dev storage_image_ID

(Optional). Specifies the storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified logical subsystem (LSS) ID.

Example: IBM.2107-75FA120

-metrics

(Optional). Specifies that the logical subsystem ID and its performance statistics 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 (2107 only) .

If you specify the dash (-), this parameter information is automatically supplied.

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

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. A separate example is not shown for the 1750 as the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

Invoking the showgmir command

```
dscli>showgmir  
-dev IBM.2107-75FA120 10
```

The resulting output

Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

ID	Master Count	Master Session ID	Copy State	Fatal Reason	CG Interval	Global Mirror Interval	CG Drain Time	Current Time
IBM. 2107- 75FA120 /10	1	01	Running	00	60	50	240	01/03 /2005 08:00:00 MST 0

CG Time	CG Failed Attempt	Successful CG Percentage	Flash-Copy Sequence Number	Master ID	Subordinate Count	Master SSID	Subordinate SSID
01/03 /2005 07:58:00 MST 0	0	100	10	IBM. 2107- 75FA120	1	0010	0020

Report field definitions

ID Specifies the LSS ID which consists of a *Storage Image ID* followed by two hexadecimal characters 00 - FE (2107 only) that identify a Global Mirror (Asynchronous PPRC) master LSS ID.

Master Count

Specifies the number of Global Mirror (Asynchronous PPRC) masters. This value could be null (-) if none exist.

Master Session ID

Identifies the session ID that you assigned, '01-FF' (1-255 decimal).

Copy State

Identifies the Global Mirror (Asynchronous PPRC) copy state. The value displayed is one of the following

Running

Running indicates that the remote copy and mirror copy process is running.

Paused

Paused indicates that the remote copy and mirror copy process is paused.

Fatal

Fatal indicates that the remote copy and mirror copy process is failed.

Pause In Progress

Pause In Progress indicates that the remote copy and mirror copy process is in the process of pausing.

Fatal Reason

Specifies a reason code for the failure.

CG Interval

Indicates the interval time between attempts to form a consistency group, up to 65 535 seconds.

Global Mirror Interval

Indicates the extended distance consistency maximum coordination interval, up to 65 535 milliseconds.

CG Drain Time

The CG 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 value 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: Daylight Savings is set to 1 if the clock is to be automatically adjusted at applicable times between standard and daylight savings times. Set to 0 if the clock is not to be automatically adjusted for daylight savings time.

CG Time

Indicates the time recorded when the last successful Consistency Group was formed.

CG Failed Attempt

Indicates the number of unsuccessful attempts to form a Consistency Group since the last time a Consistency Group was formed.

Successful CG Percentage

Indicates the percentage of attempts to form a Consistency Group that was successful, from 0% to 100%.

FlashCopy Sequence Number

Specifies the FlashCopy sequence number associated with the current Consistency Group.

Note: This value does not apply to a 2105 so a null (-) is displayed when a 2105 is associated with this reporting.

Master ID

Identifies the Global Mirror (Asynchronous PPRC) master *Storage_Image_ID*.

Subordinate Count

Specifies the count of subordinate associations. The master-subordinate association fields repeat according this count.

Master SSID

Identifies the SSID value that is associated with the master LSS ID.

Subordinate SSID

Identifies the SSID value that is associated with the subordinate LSS ID. This field repeats for each master-subordinate path association.

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. A separate example is not shown for the 1750 as the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

Invoking the showgmir command

```
dscli>showgmir -dev IBM.2107-75FA120 -metrics 10
```

The resulting output

Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

ID	Total Failed CG Count	Total Successful CG Count	Successful CG Percentage	Failed CG after Last Success	Last Successful CG Form Time	Coord. Time (milliseconds)	Interval Time (seconds)
IBM.2107-75FA120/10	2	3	40	2	01/13/1970 13:08:37 PST	50	5

Max Drain Time (seconds)	First Failure Control Unit	First Failure LSS	First Failure Status	First Failure Reason	First Failure Master State	Last Failure Control Unit	Last Failure LSS
240	IBM.2107-75FA120	0x05	Error	Long Busy	Error Recovery	IBM.2107-75FA120	0x05

Last Failure Status	Last Failure Reason	Last Failure Master State	Previous Failure Control Unit	Previous Failure LSS	Previous Failure Status	Previous Failure Reason	Previous Failure Master State
Error	Long Busy	Error Recovery	IBM.2107-75FA120	0x05	Error	Long Busy	Error Recovery

Report field definitions

ID Identifies the LSS ID. This ID consists of an *Storage_Image_ID* followed by two hexadecimal characters 00-FE (2107 only) that identify a Global Mirror (Asynchronous PPRC) master LSS ID.

Total Failed CG Count

Specifies the total number of Consistency Groups that did not complete in the user specified drain time.

Total Successful CG Count

Identifies the total number of Consistency Groups that completed before the user specified drain time.

Successful CG Percentage

Identifies the percentage of attempts that were successful in forming a consistency group.

Failed CG after Last Success

Specifies the total number of failed Consistency Groups after the last successful completion.

Last Successful CG Form Time

Identifies the last successful Consistency Group completion time.

Coord. Time (milliseconds)

Specifies the value in milliseconds that indicates the maximum amount of time Global Mirror will queue Primary/Host I/O to start forming a consistency group. Default is 50 milliseconds.

Interval Time (seconds)

Specifies the value in seconds that indicates how long to wait between formation of Consistency groups.

Max Drain Time (seconds)

Specifies the value in seconds that indicates the maximum amount of time that Global Mirror will allow for the consistent set of data to drain to the remote site.

First Failure Control Unit

Identifies the Control Unit MTS that has caused the first failure of the consistency group formation.

First Failure LSS

Identifies 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

Specifies the error reason of the first failure of the consistency group formation attempt.

First Failure Master State

Identifies the master state for the first Global Mirror failure.

Last Failure Control Unit

Identifies the Control Unit MTS that has caused the last failure of the consistency group formation.

Last Failure LSS

Identifies 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

Specifies the error reason of the last failure of the consistency group formation attempt.

Last Failure Master State

Identifies the master state for the last Global Mirror failure.

Previous Failure Control Unit

Identifies the Control Unit MTS that has caused the previous failure of the consistency group formation.

Previous Failure LSS

Identifies 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

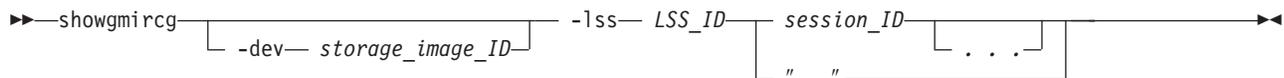
Specifies the error reason of the previous failure of the consistency group formation attempt.

Previous Failure Master State

Specifies the master state for the previous Global Mirror failure.

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, type, and serial number. This parameter is required if you do not specify a fully qualified LSS ID, which includes a storage image ID.

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 (2107 only) .

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.

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

session_ID . . . | -

(Required). Specifies one or more sessions to display. A session ID is a Global Mirror session number that you assign in the 01 - FF hexadecimal range or the 1 - 255 decimal range.

If you specify the dash (-), this parameter information is automatically supplied.

Example: 01

Example

A separate example is not shown for the 1750 as the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

Invoking the showgmircg command

```
dsccli>showgmircg -dev IBM.2107-75FA120 -lss 01 10
```

The resulting output

LSS ID IBM.2107-75FA120/10
Session 01
CG Status 0

Report field definitions

The following supplies information that you can use to help you understand the output that is generated from this command.

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— 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, type, and serial number. This parameter is required if you do not specify a fully qualified LSS ID, which includes a storage image ID.

Example: IBM.2107-75FA120

-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 (2107 only) .

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

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 or the 1 - 255 decimal range.

If you specify the dash (-), this parameter information is automatically supplied.

Example

A separate example is not shown for the 1750 as the information is the same for both. The only difference is the machine type designation, 2107 versus 1750.

Invoking the showgmirroos command

```
dscli>showgmirroos -dev IBM.2107-75FA120 -scope si -lss 10 01
```

The resulting output

```
Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0 DS: IBM.2107-75FA120
```

```
Scope IBM.2107-75FA120
Session 01
OutOfSyncTracks 3
```

Report field definitions

The following supplies information that you can use to help you understand the output that is generated from this command.

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

This section contains commands that are used to create, modify, and delete Global Mirror sessions and to display Global Mirror session information.

Use the following commands to create, modify, and delete Global Mirror sessions and to display Global Mirror session information:

- **mksession**
- **chsession**
- **lssession**
- **rmsession**

The **mksession** command opens a Global Mirror session.

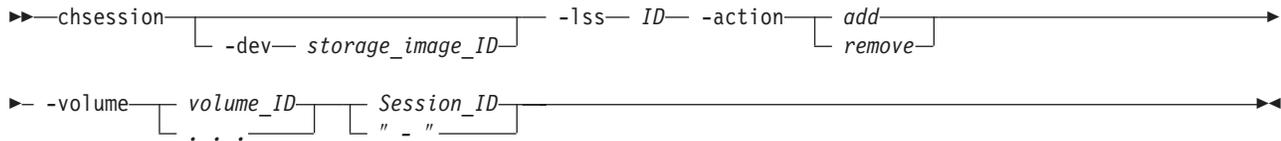
The **chsession** command allows you to modify a Global Mirror session.

The **lssession** command 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.

The **rmssession** command 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, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the LSS ID and you do not specify a value for the *devid* variable in your profile file.

Example: IBM.2107-75FA120

-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 (2107 only).

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.

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

-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.

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 4 hexadecimal digits in the form of XYZZ where X is the address group (0 - 1), XY together is the logical subsystem number 00 - FE (for 2107), and ZZ is 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

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 specify the dash (-), this parameter information is automatically supplied.

Example of a session ID: 01

Example (2107)

Invoking the `chsession` command

```
dscli>chsession -dev IBM.2107-75FA120 -lss 10 -action add  
-volume 1000-1010 01
```

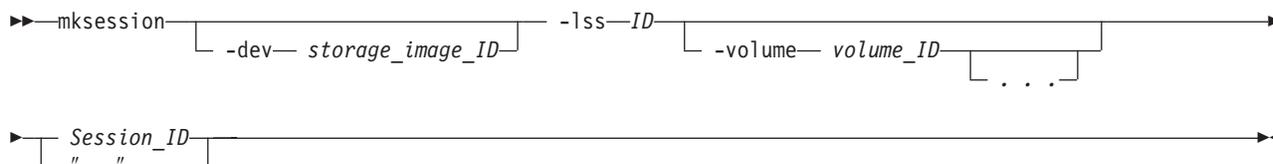
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

Global Mirror session 01 successfully modified.

mksession

The `mksession` command opens a Global Mirror session.



Parameters

-dev *storage_image_ID*

(Optional). Specifies the storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the logical subsystem.

Example: IBM.2107-75FA120

-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 (2107 only) .

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

-volume *volume_ID . . .*

(Optional). 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.

Do not qualify the volume ID with the storage image ID. The volume ID is a 32 bit number that can be represented as 4 hexadecimal digits in the form of

XYZZ where X is the address group (0 - 1), XY together is the logical subsystem number 00 - FE (for 2107), and ZZ is 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

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 specify the dash (-), this parameter information is automatically supplied.

Example of a session ID: 01

Example (2107)

Invoking the `mksession` command

```
dscli>mksession -dev IBM.2107-75FA120 -lss 10 -volume 1000-100F 01
```

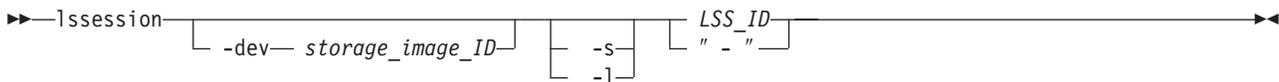
The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

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 regarding the volumes of each session in the list.



Parameters

`-dev storage_image_ID`

(Optional). Specifies the storage image ID, which includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified ID for the logical subsystem.

Example: IBM.2107-75FA120

`-s` (Optional). Displays the session IDs. You cannot use the `-l` and the `-s` parameters together.

`-l` (Optional). Displays the default output. You cannot use the `-l` and the `-s` parameters together.

LSS_ID | -

(Required). Specifies the logical subsystem (LSS) ID for the Global Mirror session. The format of the LSS ID is a hexadecimal value in the range 00 - FE (2107 only).

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 specify the dash (-), this parameter information is automatically supplied.

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

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 using the **-l** parameter. A separate example is not shown for the 1750 because the information is the same for both machine types. The only difference is the machine type designation, 2107 versus 1750.

When you use the **-s** parameter with the **lssession** command only 3 ID items are displayed in the resulting report: LSSID, SessionID, and VolumeID. A separate example is shown for this scenario.

Invoking the lssession command using the -l parameter

```
dscli>lssession -dev IBM.2107-75FA120 -l 01
```

The resulting output

Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120

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 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 Simplex	Secondary Simplex	True	Disabled
Primary Simplex	Secondary Simplex	False	Disabled

Invoking the lssession command using the -s parameter

```
dscli>lssession -s -dev IBM.2107-75FA120 10
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 IBM DS CLI Version: 5.0.0.0
DS: IBM.2107-75FA120
```

LSSID	Session	Volume
10	01	1001
10	01	1002
10	01	1003
10	02	1011
10	02	1012
10	02	1013

Report field definitions

LSSID

Specifies the unique identifier that is assigned to this logical subsystem object. The LSS ID is a hexadecimal value in the range 00 - FE (2107 only) .

Session

Specifies the Session ID number that you assigned in the 01 - FF hexadecimal range or the 1 - 255 decimal range.

Status Specifies the state of the session. One of the following values is displayed:

CG in progress

Indicates that the consistency group of the session is in progress.

Increment Pending

Indicates that the Increment process is in progress.

Normal

Indicates that the session is in normal state.

Volume

Specifies the volume ID. If no volume is active for the session a null (-) value is displayed.

Volume Status

Specifies 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 will be added to the session in the next cycle.

Remove Pending

Indicates that the volume is active for the current session. However, it will be removed in the next cycle.

Active Indicates that the volume is an active member of the session.

Primary Status

Specifies 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.

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.

Primary suspended

Indicates that the volume is primary in a remote mirror and copy relationship and the relationship is suspended.

- (null)

Indicates that there are no active volumes for the session.

Secondary Status

Specifies the primary 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.

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.

Secondary suspended

Indicates that the volume is secondary in a remote mirror and copy relationship and the relationship is suspended.

- (null)

Indicates that there are no active volumes for the session.

FirstPass Complete

Specifies whether the first cycle of the volume in the global mirror relationship has ended. The value displayed is either True or False.

AllowCascading

Specifies whether the volume can be a secondary in a remote mirror and copy relationship. The value displayed is either Enabled or Disabled.

rmssession

The **rmssession** command closes an existing Global Mirror session.

```
rmssession [-dev storage_image_ID] -lss ID [-quiet] [Session_ID]
```

Parameters

-dev storage_image_ID

(Optional). Specifies the ID of the storage image that contains the logical subsystem. The storage image ID includes manufacturer, type, and serial number. This parameter is required if you do not specify a fully qualified LSS ID.

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 (2107 only) .

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.

Example of a fully qualified LSS ID: IBM.2107-75FA120/10

-quiet

(Optional). Turns off the 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 or a decimal number in the range 01 - 255 that you assigned.

If you specify a dash (-), this parameter information is automatically supplied.

Example of a session ID: 01

Example (2107)

Invoking the rmssession command

```
dscli>rmssession -dev IBM.2107-75FA120 -lss 10 01
```

The resulting output

```
Date/Time: Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0  
DS: IBM.2107-75FA120
```

```
Are you sure you want to stop Global Mirror Session ID 01? y/n Y  
Global Mirror Session ID 01 successfully stopped.
```

Audit commands

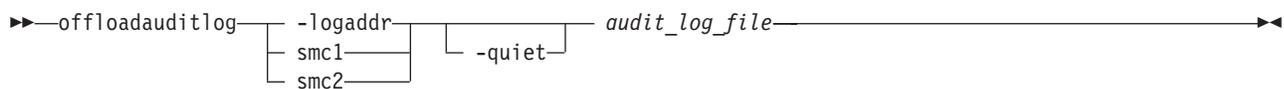
This section contains commands that are used for audit logging.

Use the **offloadauditlog** command to create activity reports for the consoles and storage units.

The **offloadauditlog** command 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.

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.



Parameters

Notes:

1. Only users with administrator authority are authorized to use this command.
2. A separate log entry is added each time a resource is created, deleted, 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.

-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) Specifies that the confirmation prompt be turned off.

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

Invoking the offloadauditlog command

```
dscli>dscli> offloadauditlog
-logaddr smc1 auditlog-200509.txt
```

The resulting output

Sun Aug 11 02:23:49 PST 2004 DS CLI Version: 5.0.0.0:
Audit log successfully offloaded from smc1 to file auditlog-200509.txt.

Representative report

The following is 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"
```

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 Parm's 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 Parm's field is enclosed in double quotation marks.

Field	Format	Description
Source	1 char	Specifies the source of the log entry: S Represents a server event that is not associated with a user action U Represents a user-requested action C Represents a continuation line for additional input attributes. There can be multiple C entries for a given user-requested (U) log entry.
Timestamp	YYYY/MM/DD HH:MM:SS:MMM TMZ	Represents the date, time, and time zone of the log entry.
User	1 - 16 char	Represents the user account that is making the request.
MC	1 char, a "1" or "2"	Represents the management console that processed the user request.
Device	16 char	Represents the storage image ID that consists of the following values: manufacture, type, and serial number.
NWC	1 char	Represents the following message types: N = notification, W = warning, and C = critical.
Entry ID	4 char	Represents the unique identifier that is associated with the activity that is represented by the log entry.
Entry name	20 char max	A text description that corresponds to the Entry ID.
Object ID	5 char max	Represents a unique identifier that identifies the object.
Exit code	8 char	Represents the final result code.

Field	Format	Description
Input Parameters	160 char max	Represents unformatted text that includes input parameters in the format: "attr1 = value1, attr2 = value2" with a comma (,) separator between parameters and double quotation marks around the entire field.

Chapter 9. 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.

The following tasks are illustrated:

- Determining the physical configuration characteristics of a storage unit.
- Determining the current logical storage and Copy Services configuration settings.
- Creating new logical storage and Copy Services configuration settings.
- Modifying or deleting logical storage and Copy Services configuration settings.

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, thereby 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, perform 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.

2. 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
```

You can add or remove volume IDs to the list in order to add or remove volumes. 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 application 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) in order to restore the physical resource to an “equivalent to new” state.

To delete fixed block data storage, perform the following steps:

1. Remove host access to the volumes that will be removed. This generally 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 -l
```

- 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.

3. 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. The shortened version of the ID is allowed if you designate the **-dev** parameter.

- 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. The shortened version of the ID is allowed if you designate the **-dev** parameter.
  - 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 -1
```

**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. The shortened version of the ID is allowed if you designate the **-dev** parameter.
- 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 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.

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, it is possible you will see a list of arrays that have a state of **unavailable**. This is generally 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) in order 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, perform the following steps:

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. 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.
2. 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 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 generally 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 only be processed using the DS CLI remote FlashCopy commands. These transactions cannot be managed through the GUI.

You must be logged into the DS CLI application 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 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.

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. Issue the **lsavailpprcport** command to obtain a report that lists which volumes are available for use. Enter the **lsavailpprcport** command at the dscli command prompt with the following parameters and variables:

```
dscli> lsavailpprcport -dev IBM.2107-1300861
-remotedev IBM.2107-1300871 -remotewwnn 5005076303FFC03D 22:2A
```

The following report is generated.

Local port	Attached port	Type
I0030	I0031	FCP
I0031	I0030	FCP
I0100	I0101	FCP
I0101	I0100	FCP

- b. Issue the **mkpprcpath** command to establish the remote mirror and copy path between LSS 22 and LSS 2A. Enter the **mkpprcpath** command at the dscli command prompt with the following parameters and variables:

```
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.

```
Date/Time: February 7, 2005 4:25:04 PM IST IBM DSCLI Version: 0.0.0.0 DS:
IBM.2107-1300861
CMUC00149I mkpprcpath: Remote Mirror and Copy path 22:2A successfully
established.
```

2. Issue the **mkpprc** command to establish a remote mirror and copy pair (2200 to 2A00). Enter the **mkpprc** command at the dscli command prompt with the following parameters and variables:

```
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.

```
Date/Time: February 7, 2005 4:25:33 PM IST IBM DSCLI Version: 0.0.0.0 DS:
IBM.2107-1300861
CMUC00153I mkpprc: Remote Mirror and Copy volume pair relationship
2200:2A00 successfully created.
```

- Issue the **mkremoteflash** command to use LSS 22 on the local site as a conduit LSS for new remote Flash Copy 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). Enter the **mkremoteflash** command at the dscli command prompt with the following parameters and variables:

```
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.

```
Date/Time: February 7, 2005 4:28:11 PM IST IBM DSCLI Version: 0.0.0.0
DS: IBM.2107-1300871
CMUC00173I mkremoteflash: Remote FlashCopy volume pair 2A00:2A01 successfully
created. Use the lsremoteflash command to determine copy completion.
```

- Issue the **resyncremoteflash** command because the remote FlashCopy relationship (2A00:2A01) was created with the **-record** parameter. Enter the **resyncremoteflash** command at the dscli command prompt with the following parameters and variables:

```
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.
```

- Issue the **lsremoteflash** command to verify that the transaction has processed as you intended. Enter the **lsremoteflash** command at the dscli command prompt with the following parameters and variables:

```
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 bring up 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 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.)

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 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:

- 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 represents an example of 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 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.
```

When this command processes, the following occurs:

- 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, perform 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 -src1ss 07 -tgt1ss 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.
```

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 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.
```

When this command processes, the following occurs:

- 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 Model 2107 storage unit 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.

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 in order 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
```

Chapter 10. 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:

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 instance, 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.

For the list of individual messages see the IBM System Storage DS8000 Information Center.

OpenVMS messages and exit codes

On OpenVMS host systems, the DS CLI application provides additional system messages in the OpenVMS message format. For example, **%facility-level-identification, text**

Each DS CLI message is also communicated by OpenVMS system services. This ensures that messages appear on the screen and in batch logs according to OpenVMS standards. This allows the OpenVMS command **SET MESSAGE** and customer-written tools which scan for such messages to work correctly. The following mapping between DS CLI and OpenVMS messages applies:

- The facility is IBMDSCLI.
- The level is I (Informational), W (Warning), or E (Error), according to the DS CLI severity level.
- The identification is the DS CLI message identifier.
- The text is a is a generic text, identifying the DS CLI message.

After displaying the OpenVMS message, the normal DS CLI message is displayed unchanged. This ensures that the DS CLI messages are identical across platforms and that you can reference the DS CLI documentation.

When you exit the DS CLI, it provides a condition code according to OpenVMS standards instead of the normal DS CLI exit value. The following rules apply:

- A *single-shot* DS CLI command returns the condition code of the DS CLI message for this command. If there are multiple messages, the code with the highest severity is displayed.
- The DS CLI in *interactive* mode always returns the OpenVMS generic success status: %X00000001.
- The DS CLI in *script* mode returns either the OpenVMS generic success code or cancels the script while returning the highest severity code of the failed command.

This ensures that the DCL symbol \$STATUS contains a value that can be evaluated by OpenVMS utilities as usual.

Accessibility

Accessibility features provide users who have disabilities with the ability to successfully access information and use technology.

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features

These are the major accessibility features in the IBM System Storage DS8000 information:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. IBM Home Page Reader version 3.0 has been tested.
- You can operate features using the keyboard instead of the mouse.

Navigating by keyboard

You can use keys or key combinations to perform operations and initiate menu actions that can also be done through mouse actions. You can navigate the IBM System Storage DS8000 information from the keyboard by using the shortcut keys for your browser or Home Page Reader. See your browser Help for a list of shortcut keys that it supports. See the following Web site for a list of shortcut keys supported by Home Page Reader: http://www-306.ibm.com/able/solution_offerings/keyshort.html

Accessing the publications

You can find HTML versions of the IBM System Storage DS8000 information at the following Web site: <http://www.ehone.ibm.com/public/applications/publications/cgibin/pbi.cgi>

You can access the information using IBM Home Page Reader 3.0.

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.*

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.

Terms and conditions

Permissions for the use of these publications is granted subject to the following terms and conditions.

Personal Use: You may reproduce these Publications for your personal, non commercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these Publications, or any portion thereof, without the express consent of IBM.

Commercial Use: You may reproduce, distribute and display these Publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these Publications, or reproduce, distribute or display these Publications or any portion thereof outside your enterprise, without the express consent of IBM.

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the Publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

- AIX
- DB2
- DFSMS/MVS
- DFSMS/VM
- DS4000
- DS6000
- DS8000
- e (logo)
- Enterprise Storage Server
- ES/9000
- ESCON
- FICON
- FlashCopy
- Graphically Dispersed Parallel Sysplex
- HACMP
- i5/OS
- IBM
- IntelliStation
- MVS/ESA
- Netfinity
- NetVista
- Operating System/400
- OS/400
- RS/6000
- S/390
- Seascape
- SNAP/SHOT
- SP
- System/390
- System p5
- System Storage
- Versatile Storage Server
- Virtualization Engine
- VSE/ESA
- z/Architecture
- z/OS
- z/VM
- zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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

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

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

Other company, product, and service names may be trademarks or service marks of others.

Electronic emission notices

This section contains the electronic emission notices or statements for the United States and other countries.

Federal Communications Commission (FCC) statement

This equipment has been tested and complies with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, might cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors, or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the users authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device might not cause harmful interference, and (2) this device must accept any interference received, including interference that might cause undesired operation.

Industry Canada compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conform à la norme NMB-003 du Canada.

European community compliance statement

This product is in conformity with the protection requirements of EC Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any

failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

Germany only

**Zulassungsbescheinigung laut Gesetz ueber die elektromagnetische
Vertraeglichkeit von Geraeten (EMVG) vom 30. August 1995.**

Dieses Geraet ist berechtigt, in Uebereinstimmung mit dem deutschen EMVG das EG-Konformitaetszeichen - CE - zu fuehren.

Der Aussteller der Konformitaetserklaeung ist die IBM Deutschland.

Informationen in Hinsicht EMVG Paragraph 3 Abs. (2) 2:

Das Geraet erfuehlt die Schutzanforderungen nach EN 50082-1 und EN 55022 Klasse A.

EN 55022 Klasse A Geraete beduerfen folgender Hinweise:

Nach dem EMVG:

"Geraete duerfen an Orten, fuer die sie nicht ausreichend entstoert sind, nur mit besonderer Genehmigung des Bundesministeriums fuer Post und Telekommunikation oder des Bundesamtes fuer Post und Telekommunikation betrieben werden. Die Genehmigung wird erteilt, wenn keine elektromagnetischen Stoerungen zu erwarten sind." (Auszug aus dem EMVG, Paragraph 3, Abs.4)

Dieses Genehmigungsverfahren ist nach Paragraph 9 EMVG in Verbindung mit der entsprechenden Kostenverordnung (Amtsblatt 14/93) kostenpflichtig.

Nach der EN 55022:

"Dies ist eine Einrichtung der Klasse A. Diese Einrichtung kann im Wohnbereich Funkstoerungen verursachen; in diesem Fall kann vom Betreiber verlangt werden, angemessene Massnahmen durchzufuehren und dafuer aufzukommen."

Anmerkung:

Um die Einhaltung des EMVG sicherzustellen, sind die Geraete wie in den Handbuechern angegeben zu installieren und zu betreiben.

Japanese Voluntary Control Council for Interference (VCCI) class A statement

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス A 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Korean Ministry of Information and Communication (MIC) statement

Please note that this device has been certified for business use with regard to electromagnetic interference. If you find this is not suitable for your use, you may exchange it for one of residential use.

Taiwan class A compliance statement

警告使用者:

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

VS07171L

Taiwan Contact Information

This topic contains the product service contact information for Taiwan.

IBM Taiwan Product Service Contact Information:
IBM Taiwan Corporation
3F, No 7, Song Ren Rd., Taipei Taiwan
Tel: 0800-016-888

台灣IBM 產品服務聯絡方式：
台灣國際商業機器股份有限公司
台北市松仁路7號3樓
電話：0800-016-888

f2c00790

Glossary

This glossary includes terms for the IBM System Storage and other Resiliency Family products.

This glossary includes selected terms and definitions from:

- *The American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI), 11 West 42nd Street, New York, New York 10036. Definitions derived from this book have the symbol (A) after the definition.
- *IBM Terminology*, which is available online at the following Web site: <http://www-306.ibm.com/ibm/terminology/index.html>. Definitions derived from this source have the symbol (GC) after the definition.
- *The Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions derived from this book have the symbol (I) after the definition. Definitions taken from draft international standards, committee drafts, and working papers that the ISO/IEC JTC1/SC1 is developing have the symbol (T) after the definition, indicating that final agreement has not been reached among the participating National Bodies of SC1.

This glossary uses the following cross-reference forms:

- See** Refers the reader to one of two kinds of related information:
- A term that is the expanded form of an abbreviation or acronym. This expanded form of the term contains the full definition.
 - A synonym or more preferred term

See also Refers the reader to one or more related terms.

Contrast with

Refers the reader to a term that has an opposite or substantively different meaning.

Numerics

- 750** A model of the Enterprise Storage Server featuring a 2-way processor with limited physical storage capacity. This model can be updated to the model 800.
- 800** A model of the Enterprise Storage Server featuring a standard processor or an optional Turbo processor. The Model 800 supports RAID 5, RAID 10, and 15000 rpm drives. Model 800 supersedes Model F20.
- 1750** The machine type for the IBM System Storage DS6000 series. Models for the DS6000 include the 511 and EX1.
- 2105** The machine number for the IBM TotalStorage Enterprise Storage Server. Models of the Enterprise Storage Server are expressed as the number 2105 followed by "Model <xxx>", such as 2105 Model 800. The 2105 Model 100 is an Enterprise Storage Server expansion enclosure that is typically referred to simply as the Model 100.
- 2107** A hardware machine type for the IBM System Storage DS8000 series. Hardware models for the 2107 include base units 921, 922, 931, 932, 9A2, 9B2 and expansion units 92E and 9AE.
- 2244** A function authorization machine type for the IBM System Storage DS8000 series. The 2244 function authorization machine type corresponds with the 2107 hardware machine type and is used only for purposes of billing and authorizing the licensed functions on the 2107. Function authorization models for the 2244 are related to the type of licensed functions that you order. For example, Model RMC is for the remote mirror and copy function on a 2107 storage unit.
- 239x** Function authorization machine types for the IBM System Storage DS8000 series. These machine types indicate the

warranty period for the licensed functions and they include the following machine types: 2396 (one-year warranty), 2397 (two-year warranty), 2398 (3-year warranty), and 2399 (four-year warranty). Each 239x function authorization machine type corresponds to the 242x hardware machine type that represents the same warranty period. For example, you order a 2398 (3-year warranty) function authorization machine type for a 2423 (3-year warranty) hardware machine. The 239x machine types are used only for purposes of billing and authorizing the licensed functions on the 242x machines. The 239x machine types have one model (Model LFA) with several types of available licenses for that model. For example, Model LFA, RMC license is for the remote mirror and copy function on a 242x storage unit.

242x Hardware machine types for the IBM System Storage DS8000 series. The 242x hardware machine types include machine types 2421 (one-year warranty), 2422 (two-year warranty), 2423 (3-year warranty), and 2424 (four-year warranty). Hardware models for the 242x machine types include base units 931, 932, 9B2 and expansion units 92E and 9AE.

3390 The machine number of an IBM disk storage system. The Enterprise Storage Server, when interfaced to IBM zSeries hosts, is set up to appear as one or more 3390 devices, with a choice of 3390-2, 3390-3, or 3390-9 track formats.

3990 The machine number of an IBM control unit.

7133 The machine number of an IBM disk storage system. The Model D40 and 020 drawers of the 7133 can be installed in the 2105-100 expansion enclosure of the ESS.

A

access 1) To obtain computing services or data.
2) In computer security, a specific type of interaction between a subject and an object that results in flow of information from one to the other.

access-any mode

One of the two access modes that can be set for the storage unit during initial

configuration. It enables all fibre-channel-attached host systems with no defined access profile to access all logical volumes on the storage unit. With a profile defined in DS Storage Manager for a particular host, that host has access only to volumes that are assigned to the WWPN for that host. See also *pseudo host* and *worldwide port name*.

ACK See *request for acknowledgment and acknowledgment*.

agent A program that automatically performs some service without user intervention or on a regular schedule. See also *subagent*.

alert A message or log that a storage unit generates as the result of error event collection and analysis. An alert indicates that a service action is required.

allegiance

For zSeries, a relationship that is created between a device and one or more channel paths during the processing of certain conditions. See also *implicit allegiance*, *contingent allegiance*, and *reserved allegiance*.

allocated storage

The space that is allocated to volumes but not yet assigned. Contrast with *assigned storage*.

American National Standards Institute (ANSI)

An organization of producers, consumers, and general interest groups that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States. (A)

anonymous

In the DS Storage Manager, the label on an icon that represents all connections that are using fibre-channel adapters between the storage unit and hosts but are not completely defined to the storage unit. See also *anonymous host*, *pseudo host*, and *access-any mode*.

anonymous host

Synonym for *pseudo host*. Contrast with *anonymous* and *pseudo host*.

ANSI See *American National Standards Institute*.

APAR See *authorized program analysis report*. (GC)

API See *application programming interface*.

application programming interface

An interface that allows an application program that is written in a high-level language to use specific data or functions of the operating system or another program.

arbitrated loop

A fibre-channel topology that enables the interconnection of a set of nodes. See also *point-to-point connection* and *switched fabric*.

array An ordered collection, or group, of physical devices (disk drive modules) that is used to define logical volumes or devices. In the storage unit, an array is a group of disks that the user designates to be managed by the RAID technique. See also *redundant array of independent disks*.

ASCII (American National Standard Code for Information Interchange) The standard code, using a coded character set consisting of 7-bit coded characters (8 bits including parity check), that is used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. (A) Some organizations, including IBM, have used the parity bit to expand the basic code set.

assigned storage

The space that is allocated to a volume and that is assigned to a port.

authorized program analysis report (APAR)

A request for correction of a defect in a current release of an IBM-supplied program. (GC)

availability

The degree to which a system or resource is capable of performing its normal function. See *data availability*.

B

bay The physical space that is used for installing SCSI, ESCON, and fibre-channel host adapter cards. The DS8000 has four bays, two in each cluster. See also *service boundary*.

bit The smallest unit of computer information, which has two possible states that are represented by the binary digits 0 or 1. See also *byte*.

block A string of data elements recorded or transmitted as a unit. The elements may be characters, words, or physical records. (GC)

A group of consecutive bytes used as the basic storage unit in fixed-block architecture (FBA). All blocks on the storage device are the same size (fixed size). See also *fixed-block architecture* and *data record*.

byte A string that represents a character and usually consists of eight binary digits that are treated as a unit. A byte is the smallest unit of storage that can be addressed directly. (GC) See also *bit*.

C

cache A special-purpose buffer storage, smaller and faster than main storage, used to hold a copy of instructions and data obtained from main storage and likely to be needed next by the processor. (GC)

cache fast write

A form of the fast-write operation in which the storage server writes the data directly to cache, where it is available for later destaging.

cache hit

An event that occurs when a read operation is sent to the cluster, and the requested data is found in cache. Contrast with *cache miss*.

cache memory

Memory, typically volatile memory, that a storage server uses to improve access times to instructions or data. The cache memory is typically smaller and faster than the primary memory or storage medium. In addition to residing in cache memory, the same data also resides on the storage devices in the storage unit.

cache miss

An event that occurs when a read operation is sent to the cluster, but the data is not found in cache. Contrast with *cache hit*.

call home

A communication link established between the storage product and a service provider. The storage product can use this link to place a call to IBM or to another service provider when it requires service.

With access to the machine, service personnel can perform service tasks, such as viewing error logs and problem logs or initiating trace and dump retrievals. (GC) See also *heartbeat* and *remote technical assistance information network*.

cascading

1) Connecting network controllers to each other in a succession of levels to concentrate many more lines than a single level permits.

2) In high-availability cluster multiprocessing (HACMP), pertaining to a cluster configuration in which the cluster node with the highest priority for a particular resource acquires the resource if the primary node fails. The cluster node relinquishes the resource to the primary node upon reintegration of the primary node into the cluster.

catcher

A server that service personnel use to collect and retain status data that an DS8000 sends to it.

CCR See *channel command retry*.

CCW See *channel command word*.

CD See *compact disc*.

central electronics complex

The set of hardware facilities that are associated with a host computer.

channel

The part of a channel subsystem that manages a single I/O interface between a channel subsystem and a set of control units.

channel command retry (CCR)

The protocol used between a channel and a control unit that enables the control unit to request that the channel reissue the current command.

channel command word (CCW)

A data structure that specifies an I/O operation to the channel subsystem.

channel path

The interconnection between a channel and its associated control units.

channel subsystem

The part of a host computer that manages I/O communication between the program and any attached control units.

channel-subsystem image

In mainframe computing, the logical functions that a system requires to perform the function of a channel subsystem. With ESCON multiple image facility (EMIF), one channel subsystem image exists in the channel subsystem for each logical partition (LPAR). Each image appears to be an independent channel subsystem program, but all images share a common set of hardware facilities. (GC)

CKD See *count key data*.

CLI See *command-line interface*. See also *Copy Services command-line interface*.

cluster

1) A partition capable of performing all DS8000 functions. With two clusters in the DS8000, any operational cluster can take over the processing of a failing cluster.

cluster processor complex

The unit within a cluster that provides the management function for the DS8000. It consists of cluster processors, cluster memory, and related logic.

command-line interface (CLI)

An interface that defines a set of commands and enables a user (or a script-like language) to issue these commands by typing text in response to the command prompt (for example, DOS commands or UNIX shell commands). See also *Copy Services command-line interface*.

compact disc

An optically read disc, typically storing approximately 660 MB. CD-ROM (compact disc read-only memory) refers to the read-only format used to distribute DS8000 code and documentation.

compression

1) The process of eliminating gaps, empty fields, redundancies, and unnecessary data to shorten the length of records or blocks.

2) Any encoding that reduces the number of bits used to represent a given message or record. (GC)

concurrent copy

A facility on a storage server that enables a program to make a backup of a data set while the logical volume remains available for subsequent processing. The

data in the backup copy is frozen at the point in time that the server responds to the request.

concurrent installation of licensed internal code
Process of installing licensed internal code on a DS8000 while applications continue to run.

concurrent maintenance
Service that is performed on a unit while it is operational.

concurrent media maintenance
Service performed on a disk drive module (DDM) without losing access to the data.

configure
In storage, to define the logical and physical devices, optional features, and program products of the input/output subsystem through the user interface that the storage unit provides for this function.

consistency group
A group of volumes participating in FlashCopy relationships in a logical subsystem, across logical subsystems, or across multiple storage units that must be kept in a consistent state to ensure data integrity.

consistency group interval time
The value in seconds that indicates the length of time between the formation of consistency groups.

consistent copy
A copy of a data entity (a logical volume, for example) that contains the contents of the entire data entity at a single instant in time.

console
A user interface to a server, for example, the interface provided on a personal computer. See also *IBM System Storage Management Console*.

contingent allegiance
In mainframe computing, a relationship that is created in a control unit between a device and a channel when the channel accepts unit-check status. The allegiance causes the control unit to guarantee access; the control unit does not present the busy status to the device. The allegiance enables the channel to retrieve sense data that is associated with the

unit-check status on the channel path associated with the allegiance. (GC)

control path
The route that is established from the master storage unit to the subordinate storage unit when more than one storage unit participates in a Global Mirror session. If there is only one storage unit (the master) in the Global Mirror session, no control path is required.

control unit (CU)
1) A device that coordinates and controls the operation of one or more input/output devices, and synchronizes the operation of such devices with the operation of the system as a whole.
2) For zSeries, a storage server with ESCON or OEMI interfaces. The control unit adapts a native device interface to an I/O interface that a zSeries host system supports.

3) The portion of the storage unit that supports the attachment of emulated count key data devices over ESCON, FICON, or OEMI interfaces. See also *cluster*.

control-unit image
In mainframe computing, a logical subsystem that is accessed through an ESCON I/O interface. One or more control-unit images exist in each control unit. Each image appears as an independent control unit, but all control-unit images share a common set of hardware facilities. The DS8000 can emulate 3990-3, TPF, 3990-6, or 2105 control units.

control-unit-initiated reconfiguration (CUIR)
A software mechanism that the DS8000 uses to request that an operating system of a zSeries host verify that one or more subsystem resources can be taken offline for service. The DS8000 can use this process to automatically vary channel paths offline and online to facilitate bay service or concurrent code installation. Depending on the operating system, support for this process might be model dependent, might depend on the IBM TotalStorage Enterprise Storage Server Subsystem Device Driver, or might not exist.

Coordinated Universal Time (UTC)

The international standard of time that is kept by atomic clocks around the world.

Copy Services

A collection of optional software features, with a Web-browser interface, used for configuring, managing, and monitoring data-copy functions.

Copy Services CLI

See *Copy Services command-line interface*.

Copy Services domain

See *Copy Services server group*.

Copy Services client

Software that runs on each DS8000 cluster in the Copy Services server group and that performs the following functions:

- Communicates configuration, status and connectivity information to the Copy Services server
- Performs data-copy functions on behalf of the Copy Services server

Copy Services command-line interface (Copy Services CLI)

The command-line interface software that is provided with DS8000 Copy Services and used for invoking Copy Services functions from host systems attached to the DS8000. See also *command-line interface*.

Copy Services server

A cluster that the Copy Services administrator designates to perform the DS8000 Copy Services functions.

Copy Services server group

A collection of user-designated DS8000 clusters participating in Copy Services functions that a designated, active, Copy Services server manages. A Copy Services server group is also called a Copy Services domain. See also *domain*.

count field

The first field of a count key data (CKD) record. This eight-byte field contains a four-byte track address (CCHH). It defines the cylinder and head that are associated with the track, and a one-byte record number (R) that identifies the record on the track. It defines a one-byte key length that specifies the length of the record's key field (0 means no key field). It defines a two-byte data length that

specifies the length of the record's data field (0 means no data field). Only the end-of-file record has a data length of zero.

count key data (CKD)

In mainframe computing, a data-record format employing self-defining record formats in which each record is represented by up to three fields: a *count* field that identifies the record and specifies its format, an optional *key* field that identifies the data area contents, and an optional *data* field that typically contains the user data. For CKD records on the storage unit, the logical volume size is defined in terms of the device emulation mode (3390 or 3380 track format). The count field is always 8 bytes long and contains the lengths of the key and data fields, the key field has a length of 0 to 255 bytes, and the data field has a length of 0 to 65 535 or the maximum that will fit on the track. See also *data record*.

CPC See *cluster processor complex*.

CRC See *cyclic redundancy check*.

CU See *control unit*.

CUIR See *control-unit initiated reconfiguration*.

custom volume

A volume in count-key-data (CKD) format that is not a standard volume, which means that it does not necessarily present the same number of cylinders and capacity to its assigned logical control unit as provided by one of the following standard zSeries volume types: 3390-2, 3390-3, 3390-9, 3390-2 (3380-track mode), or 3390-3 (3380-track mode). See also *count-key-data*, *interleave*, *standard volume*, and *volume*.

CUT See *Coordinated Universal Time*.

cyclic redundancy check (CRC)

A redundancy check in which the check key is generated by a cyclic algorithm. (T)

cylinder

A unit of storage on a CKD device with a fixed number of tracks.

D

DA See *device adapter*.

daisy chain

See *serial connection*.

DASD

See *direct access storage device*.

DASD fast write (DFW)

A function of a storage server in which active write data is stored in nonvolatile cache, thus avoiding exposure to data loss.

data availability

The degree to which data is available when needed, typically measured as a percentage of time that the system would be capable of responding to any data request (for example, 99.999% available).

data compression

A technique or algorithm used to encode data such that the encoded result can be stored in less space than the original data. The original data can be recovered from the encoded result through a reverse technique or reverse algorithm. See also *compression*.

Data Facility Storage Management Subsystem (DFSMS)

An operating environment that helps automate and centralize the management of storage. To manage storage, DFSMS provides the storage administrator with control over data class, storage class, management class, storage group, and automatic class selection routine definitions.

data field

The optional third field of a count key data (CKD) record. The count field specifies the length of the data field. The data field contains data that the program writes.

data record

The basic unit of zSeries storage on a DS8000, also known as a count-key-data (CKD) record. Data records are stored on a track. The records are sequentially numbered starting with 0. The first record, R0, is typically called the track descriptor record and contains data that the operating system normally uses to manage the track. See also *count-key-data* and *fixed-block architecture*.

data set FlashCopy

For zSeries hosts, a feature of FlashCopy

that indicates how many partial volume FlashCopy relationships are active on a volume.

data sharing

The ability of multiple host systems to concurrently utilize data that they store on one or more storage devices. The storage unit enables configured storage to be accessible to any, or all, attached host systems. To use this capability, the host program must be designed to support data that it is sharing.

DDM See *disk drive module*.

DDM group

See *disk pack*.

dedicated storage

Storage within a storage unit that is configured such that a single host system has exclusive access to the storage.

demote

To remove a logical data unit from cache memory. A storage server demotes a data unit to make room for other logical data units in the cache or because the logical data unit is not valid. The storage unit must destage logical data units with active write units before they can be demoted. See also *destage*.

destage

To move data from an online or higher priority to an offline or lower priority device. The storage unit stages incoming data into cache and then destages it to disk.

device For zSeries, a disk drive.

device adapter (DA)

A physical component of the DS8000 that provides communication between the clusters and the storage devices. The DS8000 has eight device adapters that it deploys in pairs, one from each cluster. Device adapter pairing enables the DS8000 to access any disk drive from either of two paths, providing fault tolerance and enhanced availability.

device address

For zSeries, the field of an ESCON device-level frame that selects a specific device on a control-unit image.

device ID

The unique two-digit hexadecimal number that identifies the logical device.

device interface card

A physical subunit of a storage cluster that provides the communication with the attached device drive modules.

device number

For zSeries, a four-hexadecimal-character identifier, for example 13A0, that the systems administrator associates with a device to facilitate communication between the program and the host operator. The device number is associated with a subchannel.

device sparing

A subsystem function that automatically copies data from a failing device drive module to a spare device drive module. The subsystem maintains data access during the process.

DFS See *distributed file service*.

DFSMS

See *Data Facility Storage Management Subsystem*.

direct access storage device (DASD)

- 1) A mass storage medium on which a computer stores data.
- 2) A disk device.

disk cage

A container for disk drives. Each disk cage supports eight disk packs (64 disks).

disk drive

Standard term for a disk-based nonvolatile storage medium. The DS8000 uses hard disk drives as the primary nonvolatile storage media to store host data.

disk drive module (DDM)

A field replaceable unit that consists of a single disk drive and its associated packaging.

disk drive module group

See *disk pack*.

disk drive set

A specific number of identical disk drives that have the same physical capacity and rpm.

disk pack

A group of disk drive modules (DDMs) installed as a unit in a DDM bay.

disk group

A collection of 4 disk drives that are connected to the same pair of IBM Serial Storage adapters and can be used to create a RAID array. A disk group can be formatted as count key data or fixed block, and as RAID or non-RAID, or it can be left unformatted. A disk group is a logical assemblage of disk drives. Contrast with *disk pack*.

distributed file service (DFS)

A service that provides data access over IP networks.

DNS See *domain name system*.

domain

- 1) That part of a computer network in which the data processing resources are under common control.
- 2) In TCP/IP, the naming system used in hierarchical networks.
- 3) A Copy Services server group, in other words, the set of clusters the user designates to be managed by a particular Copy Services server.

domain name system (DNS)

In TCP/IP, the server program that supplies name-to-address translation by mapping domain names to internet addresses. The address of a DNS server is the internet address of the server that hosts the DNS software for the network.

dotted decimal notation

A convention used to identify IP addresses. The notation consists of four 8-bit numbers written in base 10. For example, 9.113.76.250 is an IP address that contains the octets 9, 113, 76, and 250.

drawer

A unit that contains multiple device drive modules and provides power, cooling, and related interconnection logic to make the device drive modules accessible to attached host systems.

drive 1) A peripheral device, especially one that has addressed storage media. See also *disk drive module*.

2) The mechanism used to seek, read, and write information on a storage medium.

DS8000

See *IBM System Storage DS8000*.

DS8000 Batch Configuration tool

A program that automatically configures a DS8000. The configuration is based on data that IBM service personnel enter into the program.

DS Storage Manager

See *IBM System Storage DS Storage Manager*.

duplex

1) Regarding Copy Services, the state of a volume pair after Remote Mirror and Copy has completed the copy operation and the volume pair is synchronized.

2) In general, pertaining to a communication mode in which data can be sent and received at the same time.

dynamic sparing

The ability of a storage server to move data from a failing disk drive module (DDM) to a spare DDM while maintaining storage functions.

E

E10 The predecessor of the F10 model of the Enterprise Storage Server. See also *F10*.

E20 The predecessor of the F20 model of the Enterprise Storage Server. See also *F20*.

EC See *engineering change*.

ECKD See *extended count key data*.

eight pack

See *disk pack*.

electrostatic discharge (ESD)

An undesirable discharge of static electricity that can damage equipment and degrade electrical circuitry.

emergency power off (EPO)

A means of turning off power during an emergency, usually a switch.

EMIF See *ESCON multiple image facility*.

enclosure

A unit that houses the components of a storage subsystem, such as a control unit, disk drives, and power source.

end of file

A coded character recorded on a data medium to indicate the end of the medium. On a count-key-data direct access storage device, the subsystem indicates the end of a file by including a record with a data length of zero.

engineering change (EC)

An update to a machine, part, or program.

Enterprise Systems Architecture/390 (ESA/390)

An IBM architecture for mainframe computers and peripherals. Processor systems that follow the ESA/390 architecture include the ES/9000® family. See also *z/Architecture*.

Enterprise Systems Connection (ESCON)

1) A zSeries computer peripheral interface. The I/O interface uses zSeries logical protocols over a serial interface that configures attached units to a communication fabric.

2) A set of IBM products and services that provide a dynamically connected environment within an enterprise.

EPO See *emergency power off*.

ERDS See *error-recording data set*.

error-recording data set (ERDS)

On zSeries hosts, a data set that records data-storage and data-retrieval errors. A service information message (SIM) provides the error information for the ERDS.

error recovery procedure

Procedures designed to help isolate and, where possible, to recover from errors in equipment. The procedures are often used in conjunction with programs that record information on machine malfunctions.

ESA/390

See *Enterprise Systems Architecture/390*.

ESCD See *ESCON director*.

ESCON

See *Enterprise System Connection*.

ESCON channel

A zSeries channel that supports ESCON protocols.

ESCON director (ESCD)

An I/O interface switch that allows the

interconnection of multiple ESCON interfaces in a distributed-star topology.

ESCON host systems

zSeries hosts that attach to the DS8000 with an ESCON adapter. Such host systems run on operating systems that include MVS, VSE, TPF, or versions of VM.

ESCON multiple image facility (EMIF)

For zSeries, a function that enables LPARs to share an ESCON channel path by providing each LPAR with its own channel-subsystem image.

EsconNet

In the DS Storage Manager, the label on a pseudo host icon that represents a host connection that uses the ESCON protocol and that is not completely defined on the DS8000. See also *pseudo host* and *access-any mode*.

ESD See *electrostatic discharge*.

eServer

See *IBM eServer*.

ESSNet

See *IBM TotalStorage Enterprise Storage Server Network*.

extended count key data (ECKD)

An extension of the count key data (CKD) architecture.

extent A continuous space on a disk that is occupied by or reserved for a particular data set, data space, or file. The unit of increment is a track. See also *multiple allegiance* and *parallel access volumes*.

extent pool

A groups of extents. See also *extent*.

F

fabric In fibre channel technology, a routing structure, such as a switch, receives addressed information and routes to the appropriate destination. A fabric can consist of more than one switch. When multiple fibre-channel switches are interconnected, they are said to be *cascaded*.

failback

Pertaining to a cluster recovery from failover following repair. See also *failover*.

failover

Pertaining to the process of transferring all control to a single cluster when the other cluster in the storage unit fails. See also *cluster* and *failback*.

fast write

A write operation at cache speed that does not require immediate transfer of data to a disk drive. The subsystem writes the data directly to cache, to nonvolatile storage, or to both. The data is then available for destaging. A fast-write operation reduces the time an application must wait for the I/O operation to complete.

FATA See *fibre-channel ATA*.

FBA See *fixed-block architecture*.

FC See *feature code*. **Note:** FC is a common abbreviation for fibre channel in the industry, but the DS8000 customer documentation library reserves FC for feature code.

FC-AL See *Fibre Channel ATA*.

FCP See *Fibre Channel Protocol*.

FCS See *Fibre Channel standard*.

feature code (FC)

A code that identifies a particular orderable option and that is used by service personnel to process hardware and software orders. Individual optional features are each identified by a unique feature code.

fibre channel

A data-transmission architecture based on the ANSI Fibre Channel standard, which supports full-duplex communication. The DS8000 supports data transmission over fiber-optic cable through its fibre-channel adapters. See also *Fibre Channel Protocol* and *Fibre Channel standard*.

fibre-channel ATA (FATA)

A hard drive that combines a fibre channel interface with an ATA drive. FATAs, which provide the high performance and capacity of an ATA drive, can be used wherever fibre channel drives can connect.

Fibre Channel Arbitrated Loop (FC-AL)

An implementation of the Fibre Channel Standard that uses a ring topology for the

communication fabric. Refer to American National Standards Institute (ANSI) X3T11/93-275. In this topology, two or more fibre-channel end points are interconnected through a looped interface. This topology directly connects the storage unit to an open systems host without going through a fabric switch.

Fibre Channel Connection (FICON)

A fibre-channel communications protocol that is designed for IBM mainframe computers and peripherals. It connects the storage unit to one or more S/390 hosts using a FICON S/390 channel either directly or through a FICON switch.

Fibre Channel Protocol (FCP)

A protocol used in fibre-channel communications with five layers that define how fibre-channel ports interact through their physical links to communicate with other ports.

Fibre Channel standard (FCS)

An ANSI standard for a computer peripheral interface. The I/O interface defines a protocol for communication over a serial interface that configures attached units to a communication fabric. The protocol has two layers. The IP layer defines basic interconnection protocols. The upper layer supports one or more logical protocols (for example, FCP for SCSI command protocols and SBICON for zSeries command protocols). Refer to American National Standards Institute (ANSI) X3.230-199x. See also *Fibre Channel Protocol*.

fibre-channel topology

An interconnection topology supported on fibre-channel adapters. See also *point-to-point connection, switched fabric, and arbitrated loop*.

Fibre Channel Switched Fabric (FC-SF)

An implementation of the Fibre Channel Standard that connects the storage unit to one or more open systems hosts through a fabric switch or connects one or more S/390 hosts that run LINUX on an Fibre Channel Protocol S/390 channel.

FICON

See *fibre-channel connection*.

FiconNet

In the DS Storage Manager, the label on a

pseudo host icon that represents a host connection that uses the FICON protocol and that is not completely defined on the DS8000. See also *pseudo host* and *access-any mode*.

field replaceable unit (FRU)

An assembly that is replaced in its entirety when any one of its components fails. In some cases, a field replaceable unit might contain other field replaceable units. (GC)

FIFO See *first-in-first-out*.

File Transfer Protocol (FTP)

In TCP/IP, an application protocol used to transfer files to and from host computers. See also *Transmission Control Protocol/Internet Protocol*.

firewall

A protection against unauthorized connection to a computer or a data storage system. The protection is usually in the form of software on a gateway server that grants access to users who meet authorization criteria.

first-in-first-out (FIFO)

A queuing technique in which the next item to be retrieved is the item that has been in the queue for the longest time. (A)

fixed-block architecture (FBA)

An architecture for logical devices that specifies the format of and access mechanisms for the logical data units on the device. The logical data unit is a block. All blocks on the device are the same size (fixed size). The subsystem can access them independently.

fixed-block device

An architecture for logical devices that specifies the format of the logical data units on the device. The logical data unit is a block. All blocks on the device are the same size (fixed size); the subsystem can access them independently. This is the required format of the logical data units for host systems that attach with a SCSI or fibre-channel interface. See also *fibre channel* and *small computer systems interface*.

FlashCopy

An optional feature of the DS8000 that

can make an instant copy of data, that is, a point-in-time copy of a volume.

FlashCopy relationship

A mapping of a FlashCopy source volume and a FlashCopy target volume that allows a point-in-time copy of the source volume to be copied to the target volume. FlashCopy relationships exist from the time that you initiate a FlashCopy operation until the storage unit copies all data from the source volume to the target volume or until you delete the FlashCopy relationship, if it is persistent.

FRU See *field replaceable unit*.

FTP See *File Transfer Protocol*.

full duplex

See *duplex*.

fuzzy copy

A function of the Global Copy feature wherein modifications to the primary logical volume are performed on the secondary logical volume at a later time. The original order of update is not strictly maintained. See also *Global Copy*.

G

GB See *gigabyte*.

GDPS See *Geographically Dispersed Parallel Sysplex*.

Geographically Dispersed Parallel Sysplex (GDPS)

A zSeries multisite application-availability solution.

gigabyte (GB)

A gigabyte of storage is 10^9 bytes. A gigabyte of memory is 2^{30} bytes.

Global Copy

An optional capability of the DS8000 remote mirror and copy feature that maintains a fuzzy copy of a logical volume on the same DS8000 or on another DS8000. In other words, all modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume at a later point in time. The original order of update is not strictly maintained. See also *Remote Mirror and Copy* and *Metro Mirror*.

Global Mirror

An optional capability of the remote mirror and copy feature that provides a 2-site extended distance remote copy. Data that is written by the host to the storage unit at the local site is automatically maintained at the remote site. See also *Metro Mirror* and *Remote Mirror and Copy*.

group In DS8000 documentation, a nickname for two different kinds of groups, depending on the context. See *disk pack* or *Copy Services server group*.

H

HA See *host adapter*.

HACMP

See *high availability cluster multiprocessing*.

hard disk drive (HDD)

1) A storage medium within a storage server used to maintain information that the storage server requires.

2) A mass storage medium for computers that is typically available as a fixed disk (such as the disks used in system units of personal computers or in drives that are external to a personal computer) or a removable cartridge.

hardware service manager

An option on an AS/400 or iSeries host that enables the user to display and work with system hardware resources and to debug input-output processors (IOP), input-output adapters (IOA), and devices.

HCD See *Hardware Configuration Data*.

HDA See *head disk assembly*.

HDD See *hard disk drive*.

hdisk An AIX term for storage space.

head disk assembly (HDA)

The portion of an HDD associated with the medium and the read/write head.

heartbeat

A status report sent at regular intervals from the DS8000. The service provider uses this report to monitor the health of the call home process. See also *call home*, *heartbeat call home record*, and *remote technical assistance information network*.

heartbeat call home record

Machine operating and service information sent to a service machine. These records might include such information as feature code information and product logical configuration information.

hierarchical storage management

- 1) A function in storage management software, such as Tivoli Storage Management or Data Facility Storage Management Subsystem/MVS (DFSMS/MVS), that automatically manages free space based on the policy that the storage administrator sets.
- 2) In AS/400 storage management, an automatic method to manage and distribute data between the different storage layers, such as disk units and tape library devices.

high availability cluster multiprocessing (HACMP)

Software that provides host clustering, so that a failure of one host is recovered by moving jobs to other hosts within the cluster.

high-speed loop (HSL)

A hardware connectivity architecture that links system processors to system input/output buses and other system units.

home address

A nine-byte field at the beginning of a track that contains information that identifies the physical track and its association with a cylinder.

hop Interswitch connection. A hop count is the number of connections that a particular block of data traverses between source and destination. For example, data traveling from one hub over a wire to another hub traverses one hop.

host See *host system*.

host adapter

A physical subunit of a storage server that provides the ability to attach to one or more host I/O interfaces.

host name

The Internet address of a machine in the network. The host name can be entered in the host definition as the fully qualified

domain name of the attached host system, such as *mycomputer.city.company.com*, or as the subname of the fully qualified domain name, for example, *mycomputer*. See also *host system*.

host processor

A processor that controls all or part of a user application network. In a network, the processing unit in which the data communication access method resides. See also *host system*.

host system

A computer, either of the mainframe (for example, zSeries) or of the open-systems type, that is connected to the DS8000. Hosts are connected through ESCON, FICON, or fibre-channel interfaces.

hot plug

Pertaining to the ability to add or remove a hardware facility or resource to a unit while power is on.

HSL See *high-speed loop*.

HyperPAV (IBM HyperPAV)

An optional licensed function that you can use in conjunction with the parallel access volumes (PAV) function. IBM HyperPAV associates the volumes with either an alias address or a specified base logical volume number. When a host system requests IBM HyperPAV processing and the processing is enabled, aliases on the logical subsystem are placed in an IBM HyperPAV alias access state on all logical paths with a given path group ID. IBM HyperPAV is only supported on FICON channel paths.

I

i5/OS The IBM operating system that runs the IBM i5/OS and eServer i5 server families of servers.

IBM eServer

The IBM brand name for a series of server products that are optimized for e-commerce. The products include the iSeries, pSeries, xSeries, and zSeries.

IBM product engineering (PE)

The third-level of IBM service support. Product engineering is composed of IBM engineers who have experience in supporting a product or who are knowledgeable about the product.

IBM Serial Storage adapter

A physical adapter based on the IBM Serial Storage architecture. IBM Serial Storage adapters connect disk drive modules to DS8000 clusters.

IBM System Storage

The brand name used to identify storage products from IBM, including the IBM System Storage DS8000. See also *IBM System Storage DS8000* and *IBM System Storage DS Storage Manager*.

IBM System Storage DS8000

A member of the IBM System Storage Resiliency Family of storage servers and attached storage devices (disk drive modules). The DS8000 delivers high-performance, fault-tolerant storage and management of enterprise data, affording access through multiple concurrent operating systems and communication protocols. High performance is provided by multiple symmetrical multiprocessors, integrated caching, RAID support for the disk drive modules, and disk access through a high-speed serial storage architecture interface.

IBM System Storage DS CLI

The command-line interface (CLI) that is specific to the DS8000.

IBM System Storage DS Storage Manager (DS Storage Manager)

Software with a Web-browser interface for configuring the DS8000.

IBM HyperPAV

See *HyperPAV*.

IBM TotalStorage Enterprise Storage Server Network (ESSNet)

A private network providing Web browser access to the Enterprise Storage Server. IBM installs the ESSNet software on an IBM workstation called the IBM TotalStorage ESS Master Console, supplied with the first ESS delivery.

IBM System Storage Management Console

(MC) An IBM workstation that acts as the focal point for configuration, Copy Services management, and maintenance for the DS8000. It includes a Web browser that provides links to the user interface, including the DS Storage Manager and the DS8000 Copy Services.

IBM System Storage Multipath Subsystem Device Driver (SDD)

IBM software that provides multipath configuration support for a host system that is attached to storage devices. SDD provides enhanced data availability, dynamic input/output load balancing across multiple paths, and automatic path failover protection.

IBM System Storage Resiliency Family

A set of hardware and software features and products, as well as integrated software and services that are available on the IBM System Storage DS8000 and the IBM TotalStorage Enterprise Storage Server, Models 750 and 800.

image See *storage image*.

IML See *initial microcode load*.

implicit allegiance

In Enterprise Systems Architecture/390, a relationship that a control unit creates between a device and a channel path when the device accepts a read or write operation. The control unit guarantees access to the channel program over the set of channel paths that it associates with the allegiance.

initial microcode load (IML)

The action of loading microcode for a computer into that computer's storage.

initial program load (IPL)

The action of loading software into a computer, typically an operating system that controls the computer.

initiator

A SCSI device that communicates with and controls one or more targets. Contrast with *target*.

i-node The internal structure in an AIX operating system that describes the individual files in the operating system. It contains the code, type, location, and owner of a file.

input/output (I/O)

Pertaining to (a) input, output, or both or (b) a device, process, or channel involved in data input, data output, or both.

input/output configuration data set

A configuration definition built by the I/O configuration program (IOCP) and stored on disk files associated with the processor controller.

interleave

To automatically create two striped partitions across the drives in a RAID-5 array, both of which use the count-key-data (CKD) record format.

Internet Protocol (IP)

In the Internet suite of protocols, a protocol without connections that routes data through a network or interconnecting networks and acts as an intermediary between the higher protocol layers and the physical network. The upper layer supports one or more logical protocols (for example, a SCSI-command protocol and a zSeries command protocol). Refer to ANSI X3.230-199x. The IP acronym is the IP in TCP/IP. See also *Transmission Control Protocol/Internet Protocol*.

invalidate

To remove a logical data unit from cache memory because it cannot support continued access to the logical data unit on the device. This removal might be the result of a failure within the storage server or a storage device that is associated with the device.

I/O See *input/output*.

I/O adapter (IOA)

An input-output adapter on the PCI bus.

IOCDs

See *input/output configuration data set*.

IOCP See *I/O Configuration Program*.

I/O Configuration Program (IOCP)

A program that defines to a system all the available I/O devices and channel paths.

I/O device

An addressable read and write unit, such as a disk drive device, magnetic tape device, or printer.

I/O interface

An interface that enables a host to perform read and write operations with its associated peripheral devices.

I/O Priority Queueing

A facility in the Workload Manager of zSeries that enables the system administrator to set priorities for queueing I/Os from different system images. See also *multiple allegiance* and *parallel access volumes*.

I/O processor (IOP)

Controls input-output adapters and other devices.

I/O sequential response time

The time an I/O request is queued in processor memory waiting for previous I/Os to the same volume to complete.

IP See *Internet Protocol*.

IPL See *initial program load*.

iSeries

An IBM eServer product that emphasizes integration. It is the successor to the AS/400 family of servers.

J**Java Virtual Machine (JVM)**

A software implementation of a central processing unit (CPU) that runs compiled Java code (applets and applications). (GC)

JVM See *Java Virtual Machine*.

K

KB See *kilobyte*.

key field

The second (optional) field of a count key data record. The key length is specified in the count field. The key length determines the field length. The program writes the data in the key field and uses the key field to identify or locate a given record. The subsystem does not use the key field.

kilobyte (KB)

1) For processor storage, real, and virtual storage, and channel volume, 2¹⁰ or 1024 bytes.

2) For disk storage capacity and communications volume, 1000 bytes.

Korn shell

Interactive command interpreter and a command programming language.

KPOH

See *thousands of power-on hours*.

L

LAN See *local area network*.

last-in first-out (LIFO)

A queuing technique in which the next item to be retrieved is the item most recently placed in the queue. (A)

LBA See *logical block address*.

LCU See *logical control unit*.

least recently used (LRU)

- 1) The algorithm used to identify and make available the cache space that contains the least-recently used data.
- 2) A policy for a caching algorithm that chooses to remove from cache the item that has the longest elapsed time since its last access.

LED See *light-emitting diode*.

licensed machine code

Microcode that IBM does not sell as part of a machine, but licenses to the customer. LMC is implemented in a part of storage that is not addressable by user programs. Some IBM products use it to implement functions as an alternate to hard-wired circuitry.

LIFO See *last-in first-out*.

light-emitting diode (LED)

A semiconductor chip that gives off visible or infrared light when activated.

link address

On an ESCON interface, the portion of a source or destination address in a frame that ESCON uses to route a frame through an ESCON director. ESCON associates the link address with a specific switch port that is on the ESCON director. Equivalently, it associates the link address with the channel subsystem or control unit link-level functions that are attached to the switch port.

link-level facility

The ESCON hardware and logical functions of a control unit or channel subsystem that allow communication over an ESCON write interface and an ESCON read interface.

local area network (LAN)

A computer network located on a user's premises within a limited geographic area.

local e-mail

An e-mail configuration option for storage servers that are connected to a host-system network that does not have a domain name system (DNS) server.

logical address

On an ESCON interface, the portion of a source or destination address in a frame used to select a specific channel-subsystem or control-unit image.

logical block address (LBA)

The address assigned by the DS8000 to a sector of a disk.

logical control unit (LCU)

See *control-unit image*.

logical data unit

A unit of storage that is accessible on a given device.

logical device

The facilities of a storage server (such as the DS8000) associated with the processing of I/O operations directed to a single host-accessible emulated I/O device. The associated storage is referred to as a logical volume. The logical device is mapped to one or more host-addressable units, such as a device on a zSeries I/O interface or a logical unit on a SCSI I/O interface, such that the host initiating I/O operations to the I/O-addressable unit interacts with the storage on the associated logical device.

logical partition (LPAR)

For zSeries, a set of functions that create the programming environment in which more than one logical partition (LPAR) is established on a processor. An LPAR is conceptually similar to a virtual machine environment except that the LPAR is a function of the processor. Also, the LPAR does not depend on an operating system to create the virtual machine environment. (DS8000 series only)

logical path

- 1) The relationship between a channel image and a control-unit image that designates the physical path to be used for device-level communications between these images. The logical path is established as part of the channel and control-unit initialization procedures by the exchange of link-level frames.
- 2) With the Remote Mirror and Copy feature, the relationship between a source logical subsystem (LSS) and a target LSS that is created over a physical path through the interconnection fabric that is

used for Remote Mirror and Copy functions. An LSS is a primary control unit, which performs the functions of a channel image.

logical subsystem (LSS)

A topological construct that consists of a group of up to 256 logical devices. A DS8000 can have (if CDK only) up to 32 CKD-formatted logical subsystems (8192 CKD logical devices) or (if FBA only) up to 32 fixed-block logical subsystems (8192 fixed-block logical devices). If mixed CKD and FBA, a DS8000 can have up to 16 CKD-formatted logical subsystems (4096 CKD logical devices) and up to 16 fixed-block logical subsystems (4096 fixed-block logical devices). The logical subsystem facilitates configuration of the DS8000 and might have other implications relative to the operation of certain functions. There is a one-to-one mapping between a CKD logical subsystem and a zSeries control-unit image.

For zSeries hosts, a logical subsystem represents a logical control unit (LCU). Each control-unit image is associated with only one logical subsystem. See also *control-unit image*.

logical unit

In open systems, a logical disk drive.

logical unit number (LUN)

In the SCSI protocol, a unique number that is used on a SCSI bus to enable it to differentiate between separate devices, each of which is a logical unit.

logical volume

The storage medium that is associated with a logical disk drive. A logical volume typically resides on one or more storage devices. The DS8000 administrator defines this unit of storage. The logical volume, when residing on a RAID-formatted array, is spread over the drives in the array.

logical volume manager (LVM)

A set of system commands, library routines, and other tools that allow the user to establish and control logical volume storage. The LVM maps data between the logical view of storage space and the physical disk drive module.

longitudinal redundancy check (LRC)

1) A method of error checking during data transfer that involves checking parity on a row of binary digits that are members of a set that forms a matrix. Longitudinal redundancy check is also called a longitudinal parity check.

2) A mechanism that the DS8000 uses for locating errors. The LRC checks the data as it progresses from the host, through the DS8000 controller, into the device adapter, and to the array.

longwave laser adapter

A connector that is used between a host and the DS8000 to support longwave fibre-channel communication.

loop The physical connection between a pair of device adapters in the DS8000. See also *device adapter*.

LPAR See *logical partition*.

LRC See *longitudinal redundancy check*.

LRU See *least recently used*.

LSS See *logical subsystem*.

LUN See *logical unit number*.

LVM See *logical volume manager*.

M

machine level control (MLC)

A database that contains the EC level and configuration of products in the field.

machine reported product data (MRPD)

Product data gathered by a machine and sent to a destination such as an IBM support server or RETAIN. These records might include such information as feature code information and product logical configuration information.

mainframe

A computer, usually in a computer center, with extensive capabilities and resources to which other computers may be connected so that they can share facilities. (T)

maintenance analysis procedure (MAP)

A hardware maintenance document that gives an IBM service representative a step-by-step procedure for tracing a symptom to the cause of a failure.

management console

See *IBM System Storage Management Console*.

management information base (MIB)

- 1) A collection of objects that can be accessed by means of a network management protocol. (GC)
- 2) The MIB record conforms to the Open Systems Interconnection (OSI) standard defined by the International Organization for Standardization (ISO) for the exchange of information. See also *simple network management protocol*.

MAP See *maintenance analysis procedure*.

master storage unit

The physical unit that controls the creation of consistency groups in a Global Mirror session. The master storage unit sends commands to subordinate storage units. A storage unit can be a master for only one Global Mirror session. Contrast with *subordinate storage unit*.

maximum consistency group drain time

The value in seconds that indicates the maximum time that writes from the local site are delayed to the remote site while the current consistency group is being formed at the remote site. When this time is exceeded, the current attempt to form a consistency group is ended and another attempt is started. If this time is exceeded five times, this maximum time is ignored on the next attempt to form a consistency group. The default value is the larger of four minutes or two times the consistency group interval time if this value is set to zero.

maximum coordination time

The value in milliseconds that indicates the maximum time that is allowed for host I/O to be delayed during the coordination of the primary volumes of an Global Mirror session. The default is 50 milliseconds if this value is set to zero.

MB See *megabyte*.

MC See *IBM System Storage Management Console*.

MCA See *Micro Channel architecture*.

MDM See *Multiple Device Manager*.

mean time between failures (MTBF)

- 1) A projection of the time that an individual unit remains functional. The time is based on averaging the performance, or projected performance, of a population of statistically independent units. The units operate under a set of conditions or assumptions.
- 2) For a stated period in the life of a functional unit, the mean value of the lengths of time between consecutive failures under stated conditions. (I) (A)

medium

For a storage unit, the disk surface on which data is stored.

megabyte (MB)

- 1) For processor storage, real and virtual storage, and channel volume, 2^{20} or 1 048 576 bytes.
- 2) For disk storage capacity and communications volume, 1 000 000 bytes.

Metro Mirror

A function of a storage server that maintains a consistent copy of a logical volume on the same storage server or on another storage server. All modifications that any attached host performs on the primary logical volume are also performed on the secondary logical volume. See also *Remote Mirror and Copy* and *Global Copy*.

MES See *miscellaneous equipment specification*.

MIB See *management information base*.

Micro Channel architecture (MCA)

The rules that define how subsystems and adapters use the Micro Channel bus in a computer. The architecture defines the services that each subsystem can or must provide.

Microsoft Internet Explorer

Web browser software manufactured by Microsoft.

migration

The replacement of a system or subsystem with a different type of system or subsystem, such as replacing a SCSI host adapter with a fibre-channel host adapter. In the context of data migration regarding the DS8000, the transfer of data from one storage unit to another, such as from a 3390 to the DS8000.

MIH See *missing-interrupt handler*.

mirrored pair

Two units that contain the same data. The system refers to them as one entity.

mirroring

In host systems, the process of writing the same data to two disk units within the same auxiliary storage pool at the same time.

miscellaneous equipment specification (MES)

IBM field-installed change to a machine.

missing-interrupt handler (MIH)

An MVS and MVS/XA facility that tracks I/O interrupts. MIH informs the operator and creates a record whenever an expected interrupt fails to occur before a specified elapsed time is exceeded.

MLC See *machine level control*.

mobile solutions terminal (MoST)

The mobile terminal used by service personnel.

mode conditioning patch cable

A cable that converts a single-mode signal from a longwave adapter into a light signal that is appropriate for multimode fibre. Another mode conditioning patch cable is required at the terminating end of the multimode fibre to convert the signal back to a single-mode signal for a longwave adapter.

Model 100

A 2105 Model 100, often simply referred to as a Mod 100, is an expansion enclosure for the Enterprise Storage Server. See also *2105*.

MoST See *mobile solutions terminal*.

MRPD

See *machine reported product data*.

MSA See *multiport serial adapter*.

MTBF See *mean time between failures*.

Multipath Subsystem Device Driver

See *IBM System Storage DS8000 Multipath Subsystem Device Driver*.

multiple allegiance

A DS8000 hardware function that is independent of software support. This function enables multiple system images to concurrently access the same logical volume on the DS8000 as long as the

system images are accessing different extents. See also *extent* and *parallel access volumes*.

Multiple Device Manager (MDM)

A component of the IBM TotalStorage Productivity Center that allows administrators to configure, manage, and monitor the performance of SAN storage devices from a single console.

multiple relationship FlashCopy

An option of the DS8000 that creates backup copies from one source to multiple targets by simultaneously establishing multiple FlashCopy relationships.

multiple virtual storage (MVS)

Implies MVS/390, MVS/XA, MVS/ESA, and the MVS element of the zSeries operating system.

multiplex

The action of transmitting simultaneously.

multiport serial adapter (MSA)

An adapter on the IBM System Storage Management Console that has multiple ports to which aDS8000 can be attached.

multiprocessor

A computer that includes two or more processors that have common access to a main storage. For the DS8000, the multiprocessors operate in parallel.

MVS See *multiple virtual storage*.

N

name server

A server that stores names of the participating DS8000 clusters.

near-line

A type of intermediate storage between online storage (which provides constant, rapid access to data) and offline storage (which provides infrequent data access for backup purposes or long-term storage).

Netfinity

IBM Intel-processor-based server; predecessor to the IBM xSeries server.

Netscape Navigator

Web browser software manufactured by Netscape.

network manager

A program or group of programs that is

used to monitor, manage, and diagnose the problems of a network. (GC)

node The unit that is connected in a fibre-channel network. A DS8000 is a node in a fibre-channel network.

non-RAID

A disk drive set up independently of other disk drives and not set up as part of a disk pack to store data using the redundant array of disks (RAID) data-striping methodology.

nonremovable medium

A recording medium that cannot be added to or removed from a storage device.

nonvolatile storage (NVS)

Memory that stores active write data to avoid data loss in the event of a power loss.

NVS See *nonvolatile storage*.

O

octet In Internet Protocol addressing, one of the four parts of a 32-bit integer presented in dotted decimal notation. See also *dotted decimal notation*.

OEMI See *original equipment manufacturer's information*.

open system

A system whose characteristics comply with standards made available throughout the industry and that therefore can be connected to other systems complying with the same standards. Applied to the DS8000, such systems are those hosts that connect to the DS8000 through SCSI or FCP protocols. See also *small computer system interface* and *Fibre Channel Protocol*.

organizationally unique identifier (OUI)

An IEEE-standards number that identifies an organization with a 24-bit globally unique assigned number referenced by various standards. OUI is used in the family of 802 LAN standards, such as Ethernet and Token Ring.

original equipment manufacturer's information (OEMI)

A reference to an IBM guideline for a computer peripheral interface. The interface uses ESA/390 logical protocols

over an I/O interface that configures attached units in a multidrop bus topology.

OS/390

The IBM operating system that includes and integrates functions that many IBM software products (including the MVS operating system) previously provided for the IBM S/390 family of enterprise servers.

OUI See *organizationally unique identifier*.

P

panel The formatted display of information that appears on a display screen.

parallel access volumes (PAV)

A licensed function of the DS8000 that enables OS/390 and z/OS systems to issue concurrent I/O requests against a count key data logical volume by associating multiple devices of a single control-unit image with a single logical device. Up to eight device addresses can be assigned to a PAV. The PAV function enables two or more concurrent write operations to the same logical volume, as long as the write operations are not to the same extents. See also *extent*, *I/O Priority Queueing*, and *multiple allegiance*.

parity A data checking scheme used in a computer system to ensure the integrity of the data. The RAID implementation uses parity to re-create data if a disk drive fails.

path group

In zSeries architecture, a set of channel paths that are defined to a control unit as being associated with a single logical partition (LPAR). The channel paths are in a group state and are online to the host. See also *logical partition*.

path group identifier

In zSeries architecture, the identifier that uniquely identifies a given logical partition (LPAR). The path group identifier is used in communication between the LPAR program and a device. The identifier associates the path group with one or more channel paths, thereby defining these paths to the control unit as being associated with the same LPAR. See also *logical partition*.

PAV See *parallel access volumes*.

PCI See *peripheral component interconnect*.

PDU See *protocol data unit*.

PE See *IBM product engineering*.

peripheral component interconnect (PCI)

An architecture for a system bus and associated protocols that supports attachments of adapter cards to a system backplane.

persistent FlashCopy

A state where a FlashCopy relationship remains indefinitely until the user deletes it. The relationship between the source and target volumes is maintained after a background copy completes.

physical path

A single path through the I/O interconnection fabric that attaches two units. For Copy Services, this is the path from a host adapter on one DS8000 (through cabling and switches) to a host adapter on another DS8000.

pinned data

Data that is held in cache until either an error condition is corrected and it can be moved to disk storage or until the data is discarded by a host command. Pinned data conditions can only occur on an ESS Model 800 during fast-write or dual-copy functions.

planar The main printed circuit board (PCB) that other PCBs or assemblies plug into. The planar distributes both power and signals and therefore creates a common communications path to whichever device that plugs into it.

point-in-time copy

A FlashCopy option that creates an instantaneous view of original source data at a specific moment in time.

point-to-point connection

A fibre-channel topology that enables the direct interconnection of ports. See also *arbitrated loop* and *switched fabric*.

port A physical connection on a host adapter to the cable that connects the DS8000 to hosts, switches, or another DS8000. The DS8000 uses SCSI and ESCON host adapters that have two ports per adapter, and fibre-channel host adapters that have

one port. See also *ESCON*, *fibre-channel*, *host adapter*, and *small computer system interface*.

POST See *power-on self test*.

power-on self test (POST)

A diagnostic test that servers or computers run when they are turned on.

predictable write

A write operation that can cache without knowledge of the existing format on the medium. All write operations on FBA DASD devices are predictable. On CKD DASD devices, a write operation is predictable if it does a format write operation for the first data record on the track.

primary control unit

The DS8000 to which a Remote Mirror and Copy primary device is physically attached.

processor complex

A partition of a storage server that is capable of performing all defined functions of the storage server. Multiple processor complexes provide redundancy.

product engineering

See *IBM product engineering*.

program

On a computer, a generic term for software that controls the operation of the computer. Typically, the program is a logical assemblage of software modules that perform multiple related tasks.

program-controlled interruption

An interruption that occurs when an I/O channel fetches a channel command word with the program-controlled interruption flag on.

program temporary fix (PTF)

A temporary solution to, or bypass of, a problem diagnosed by IBM as the result of a defect in a current unaltered release of a licensed program. (GC)

promote

To add a logical data unit to cache memory.

protected volume

In AS/400, a disk storage device that is protected from data loss by RAID techniques. An AS/400 host does not

mirror a volume configured as a protected volume, while it does mirror all volumes configured as unprotected volumes. The DS8000, however, can be configured to indicate that an AS/400 volume is protected or unprotected and give it RAID protection in either case.

protocol data unit (PDU)

A unit of data specified in the protocol of a given layer and consisting of protocol control information for the layer and, possibly, user data for the layer.

pSeries

The product name of an IBM eServer product that emphasizes performance. It is the successor to the RS/6000 family of servers.

pseudo host

A host connection that is not explicitly defined to the DS8000 and that has access to at least one volume that is configured on the DS8000. The FiconNet pseudo host icon represents the FICON protocol. The EsconNet pseudo host icon represents the ESCON protocol. The pseudo host icon labelled Anonymous represents hosts connected through the FCP protocol. *Anonymous host* is a commonly used synonym for *pseudo host*. The DS8000 adds a pseudo host icon only when it is set to access-any mode. See also *access-any mode*.

PTF See *program temporary fix*.

PV Links

Short for Physical Volume Links, an alternate pathing solution from Hewlett-Packard that provides for multiple paths to a volume, as well as static load balancing.

R

R0 See *track-descriptor record*.

rack See *enclosure*.

RAID See *redundant array of independent disks*. RAID is also commonly expanded to redundant array of *inexpensive* disks. See also *array*.

RAID 5

A type of RAID that optimizes cost-effective performance while emphasizing use of available capacity

through data striping. RAID 5 provides fault tolerance for up to two failed disk drives by distributing parity across all the drives in the array plus one parity disk drive. The DS8000 automatically reserves spare disk drives when it assigns arrays to a device adapter pair (DA pair). See also *device adapter*, *RAID 10*, and *redundant array of independent disks*.

RAID 10

A type of RAID that optimizes high performance while maintaining fault tolerance for up to two failed disk drives by striping volume data across several disk drives and mirroring the first set of disk drives on an identical set. The DS8000 automatically reserves spare disk drives when it assigns arrays to a device adapter pair (DA pair). See also *device adapter*, *RAID 5*, and *redundant array of independent disks*.

random access

A mode of accessing data on a medium in a manner that requires the storage device to access nonconsecutive storage locations on the medium.

rank One or more arrays that are combined to create a logically contiguous storage space.

redundant array of independent disks (RAID)

A methodology of grouping disk drives for managing disk storage to insulate data from a failing disk drive.

refresh FlashCopy target volume

An option (previously called *incremental FlashCopy*) of the DS8000 that creates a point-in-time data copy without copying an entire volume for each point-in-time copy.

Remote Mirror and Copy

A feature of a storage server that constantly updates a secondary copy of a logical volume to match changes made to a primary logical volume. The primary and secondary volumes can be on the same storage server or on separate storage servers. See also *Global Mirror*, *Metro Mirror* and *Global Copy*.

remote technical assistance information network (RETAIN)

The initial service tracking system for IBM service support, which captures

heartbeat and call-home records. See also *support catcher* and *support catcher telephone number*.

REQ/ACK

See *request for acknowledgment and acknowledgment*.

request for acknowledgment and acknowledgment (REQ/ACK)

A cycle of communication between two data transport devices for the purpose of verifying the connection, which starts with a request for acknowledgment from one of the devices and ends with an acknowledgment from the second device. The REQ and ACK signals help to provide uniform timing to support synchronous data transfer between an initiator and a target. The objective of a synchronous data transfer method is to minimize the effect of device and cable delays.

reserved allegiance

For zSeries, a relationship that is created in a control unit between a device and a channel path, or path group, when the device completes a Sense Reserve command. The allegiance causes the control unit to guarantee access (that is, busy status is not presented) to the device. Access is over the set of channel paths that are associated with the allegiance; access is for one or more channel programs until the allegiance ends.

RETAIN

See *remote technical assistance information network*.

S

S/390 IBM enterprise servers based on Enterprise Systems Architecture/390 (ESA/390). *S/390* is the currently accepted shortened form of the original name *System/390*.

S/390 storage

Storage arrays and logical volumes that are defined as connected to S/390 servers. This term is synonymous with count-key-data storage.

SAID See *system adapter identification number*.

SAM See *sequential access method*.

SAN See *storage area network*.

SBCON

See *Single-Byte Command Code Sets Connection*.

screen The physical surface of a display device upon which information is shown to users.

SCSI See *small computer system interface*.

SCSI device

A disk drive connected to a host through an I/O interface using the SCSI protocol. A SCSI device is either an initiator or a target. See also *initiator* and *small computer system interface*.

SCSI-FCP

Synonym for Fibre Channel Protocol, a protocol used to transport data between an open-systems host and a fibre-channel adapter on an DS8000. See also *Fibre Channel Protocol* and *small computer system interface*.

SCSI host systems

Host systems that are attached to the DS8000 with a SCSI interface. Such host systems run on UNIX, i5/OS, Windows NT, Windows 2000, or Novell NetWare operating systems.

SCSI ID

A unique identifier assigned to a SCSI device that is used in protocols on the SCSI interface to identify or select the device. The number of data bits on the SCSI bus determines the number of available SCSI IDs. A wide interface has 16 bits, with 16 possible IDs.

SDD See *IBM Subsystem Multipathing Device Driver*.

secondary control unit

The DS8000 to which a Remote Mirror and Copy secondary device is physically attached.

self-timed interface (STI)

An interface that has one or more conductors that transmit information serially between two interconnected units without requiring any clock signals to recover the data. The interface performs clock recovery independently on each serial data stream and uses information in

the data stream to determine character boundaries and inter-conductor synchronization.

sequential access

A mode of accessing data on a medium in a manner that requires the storage device to access consecutive storage locations on the medium.

sequential access method (SAM)

An access method for storing, deleting, or retrieving data in a continuous sequence based on the logical order of the records in the file.

serial connection

A method of device interconnection for determining interrupt priority by connecting the interrupt sources serially.

server A host that provides certain services to other hosts that are referred to as clients.

A functional unit that provides services to one or more clients over a network. (GC)

service boundary

A category that identifies a group of components that are unavailable for use when one of the components of the group is being serviced. Service boundaries are provided on the DS8000, for example, in each host bay and in each cluster.

service clearance

The area that is required to open the service covers and to pull out components for servicing.

service information message (SIM)

A message sent by a storage server to service personnel through an zSeries operating system.

service personnel

A generalization referring to individuals or companies authorized to service the DS8000. The terms *service provider*, *service representative*, and *IBM service support representative (SSR)* refer to types of service personnel. See also *service support representative*.

service processor

A dedicated processing unit that is used to service a storage unit.

service support representative (SSR)

Individuals or a company authorized to service the DS8000. This term also refers

to a service provider, a service representative, or an IBM service support representative (SSR). An IBM SSR installs the DS8000.

SES SCSI Enclosure Services.

session

A collection of volumes within a logical subsystem that are managed together during the creation of consistent copies of data. All volumes in a session must transfer their data successfully to the remote site before the increment can be called complete.

SFP Small form factor pluggables.

shared storage

Storage that is configured so that multiple hosts can concurrently access the storage. The storage has a uniform appearance to all hosts. The host programs that access the storage must have a common model for the information on a storage device. The programs must be designed to handle the effects of concurrent access.

shortwave laser adapter

A connector that is used between host and DS8000 to support shortwave fibre-channel communication.

SIM See *service information message*.

Simple Network Management Protocol (SNMP)

In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an application layer protocol. Information on devices managed is defined and stored in the application's Management Information Base (MIB). (GC) See also *management information base*.

simplex volume

A volume that is not part of a FlashCopy, XRC, or PPRC volume pair.

Single-Byte Command Code Sets Connection (SBCON)

The ANSI standard for the ESCON I/O interface.

small computer system interface (SCSI)

A standard hardware interface that enables a variety of peripheral devices to communicate with one another. (GC)

smart relay host

A mail relay or mail gateway that has the capability to correct e-mail addressing problems.

SMIT See *System Management Interface Tool*.

SMP See *symmetrical multiprocessor*.

SNMP

See *Simple Network Management Protocol*.

SNMP agent

A server process that resides on a network node and is responsible for communicating with managers regarding that node. The node is represented as a managed object, which has various fields or variables that are defined in the appropriate MIB.

SNMP manager

A managing system that runs a managing application or suite of applications. These applications depend on Management Information Base (MIB) objects for information that resides on the managed system. Managers generate requests for this MIB information, and an SNMP agent on the managed system responds to these requests. A request can either be the retrieval or modification of MIB information.

software transparency

Criteria applied to a processing environment that states that changes do not require modifications to the host software in order to continue to provide an existing function.

source device

One of the devices in a dual-copy or remote-copy volume pair. All channel commands to the logical volume are directed to the source device. The data on the source device is duplicated on the target device. See also *target device*.

spare A disk drive on the DS8000 that can replace a failed disk drive. A spare can be predesignated to allow automatic dynamic sparing. Any data preexisting on a disk drive that is invoked as a spare is destroyed by the dynamic sparing copy process.

spatial reuse

A feature of serial storage architecture that enables a device adapter loop to

support many simultaneous read/write operations. See also *serial storage architecture*.

SSID See *subsystem identifier*.

SSR See *service support representative*.

stacked status

For zSeries, the condition when the control unit is in a holding status for the channel, and the last time the control unit attempted to present the status, the channel responded with the stack-status control.

stage operation

The operation of reading data from the physical disk drive into the cache.

staging

To move data from an offline or low-priority device back to an online or higher priority device, usually on demand of the system or on request of the user.

standard volume

A volume that emulates one of several zSeries volume types, including 3390-2, 3390-3, 3390-9, 3390-2 (3380-track mode), or 3390-3 (3380-track mode), by presenting the same number of cylinders and capacity to the host as provided by the native zSeries volume type of the same name.

STI See *self-timed interface*.

storage area network

A network that connects a company's heterogeneous storage resources.

storage capacity

The amount of data that a storage medium can hold; usually expressed in kilobytes, megabytes, or gigabytes.

storage complex

A configuration of one or more storage units that is managed by a management console.

storage device

A physical unit that provides a mechanism to store data on a given medium such that it can be subsequently retrieved. See also *disk drive module*.

storage extent

The minimum contiguous range of

storage on a physical storage device, array, or rank that can be allocated to a local volume

storage image

A partitioning of a storage unit that provides emulation of a storage server with one or more storage devices that provides storage capability to a host computer. You can configure more than one storage image on a storage unit. (DS8000 series only)

storage server

A physical unit that manages attached storage devices and provides an interface between them and a host computer by providing the function of one or more logical subsystems. The storage server can provide functions that the storage device does not provide. The storage server has one or more clusters.

storage unit

A physical unit that consists of a storage server that is integrated with one or more storage devices that provide storage capability to a host computer.

storage unit identifier

A unique identifier for a storage unit that consists of a manufacturer, a model number, a type number, a plant of manufacture, and a sequence number.

striping

A technique that distributes data in bit, byte, multibyte, record, or block increments across multiple disk drives.

subagent

An extension to an SNMP agent that permits a user to dynamically add, or in some cases replace, additional management variables in the local MIB, thereby providing a means of extending the range of information that network managers can access. See also *agent*.

subchannel

A logical function of a channel subsystem associated with the management of a single device.

subordinate storage unit

The physical unit that receives commands from the master storage unit and is specified when a Global Mirror session is started. The subordinate storage unit forms consistency groups and performs

other Global Mirror processing. A subordinate storage unit can be controlled by only one master storage unit. Contrast with *master storage unit*.

subsystem identifier (SSID)

A number that uniquely identifies a logical subsystem within a computer installation.

support catcher

See *catcher*.

support catcher telephone number

The telephone number that connects the support catcher server to the DS8000 to receive a trace or dump package. See also *support catcher* and *remote technical assistance information network*.

switched fabric

A fibre-channel topology in which ports are interconnected through a switch. Fabric switches can also be interconnected to support numerous ports on a single network. See also *arbitrated loop* and *point-to-point connection*.

symmetrical multiprocessor (SMP)

An implementation of a multiprocessor computer consisting of several identical processors configured in a way that any subset of the set of processors is capable of continuing the operation of the computer. The DS8000 contains four processors set up in SMP mode.

synchronous write

A write operation whose completion is indicated after the data has been stored on a storage device.

System/390

See *S/390*.

system adapter identification number (SAID)

The unique identification number that is automatically assigned to each DS8000 host adapter for use by Copy Services.

System Management Interface Tool (SMIT)

An interface tool of the AIX operating system for installing, maintaining, configuring, and diagnosing tasks.

System Modification Program

A program used to install software and software changes on MVS systems.

T

target A SCSI device that acts as a subordinate to an initiator and consists of a set of one or more logical units, each with an assigned logical unit number (LUN). The logical units on the target are typically I/O devices. A SCSI target is analogous to a zSeries control unit. See also *small computer system interface*.

target device

One of the devices in a dual-copy or remote-copy volume pair that contains a duplicate of the data that is on the source device. Unlike the source device, the target device might only accept a limited subset of data. See also *source device*.

TB See *terabyte*.

TCP/IP

See *Transmission Control Protocol/Internet Protocol*.

terabyte (TB)

- 1) Nominally, 1 000 000 000 000 bytes, which is accurate when speaking of bandwidth and disk storage capacity.
- 2) For DS8000 cache memory, processor storage, real and virtual storage, a terabyte refers to 2^{40} or 1 099 511 627 776 bytes.

terminal emulator

A function of the management console that allows it to emulate a terminal.

thousands of power-on hours (KPOH)

A unit of time used to measure the mean time between failures (MTBF).

time sharing option (TSO)

An operating system option that provides interactive time sharing from remote terminals.

System Storage

See *IBM System Storage*.

TPF See *transaction processing facility*.

track A unit of storage on a CKD device that can be formatted to contain a number of data records. See also *home address*, *track-descriptor record*, and *data record*.

track-descriptor record (R0)

A special record on a track that follows the home address. The control program uses it to maintain certain information

about the track. The record has a count field with a key length of zero, a data length of 8, and a record number of 0. This record is sometimes referred to as R0.

transaction processing facility (TPF)

A high-availability, high-performance IBM operating system, designed to support real-time, transaction-driven applications. The specialized architecture of TPF is intended to optimize system efficiency, reliability, and responsiveness for data communication and database processing. TPF provides real-time inquiry and updates to a large, centralized database, where message length is relatively short in both directions, and response time is generally less than three seconds. Formerly known as the Airline Control Program/Transaction Processing Facility (ACP/TPF).

Transmission Control Protocol (TCP)

A communications protocol used in the Internet and in any network that follows the Internet Engineering Task Force (IETF) standards for internetwork protocol. TCP provides a reliable host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It uses the Internet Protocol (IP) as the underlying protocol.

Transmission Control Protocol/Internet Protocol (TCP/IP)

- 1) A combination of data-transmission protocols that provide end-to-end connections between applications over interconnected networks of different types.
- 2) A suite of transport and application protocols that run over the Internet Protocol. (GC) See also *Internet Protocol* and *Transmission Control Protocol*.

transparency

See *software transparency*.

TSO See *time sharing option*.

turbo processor

A faster multiprocessor that has six processors with common access to the main storage.

U

UFS UNIX filing system.

Ultra-SCSI

An enhanced small computer system interface.

unconfigure

To delete the configuration.

unit address

For zSeries, the address associated with a device on a given control unit. On ESCON interfaces, the unit address is the same as the device address. On OEMI interfaces, the unit address specifies a control unit and device pair on the interface.

unprotected volume

An AS/400 term that indicates that the AS/400 host recognizes the volume as an unprotected device, even though the storage resides on a RAID-formatted array and is, therefore, fault tolerant by definition. The data in an unprotected volume can be mirrored. Also referred to as an *unprotected device*.

upper-layer protocol

The layer of the Internet Protocol (IP) that supports one or more logical protocols (for example, a SCSI-command protocol and an ESA/390 command protocol). Refer to ANSI X3.230-199x.

UTC See *Coordinated Universal Time*.

V

virtual machine facility

A virtual data processing machine that appears to the user to be for the exclusive use of that user, but whose functions are accomplished by sharing the resources of a shared data processing system. An alternate name for the VM/370 IBM operating system.

vital product data (VPD)

Information that uniquely defines the system, hardware, software, and microcode elements of a processing system.

VM The root name of several IBM operating systems, such as VM/XA, VM/ESA, VM/CMS, and z/VM. See also *virtual machine facility*.

volume

For zSeries, the information recorded on a single unit of recording medium. Indirectly, it can refer to the unit of recording medium itself. On a nonremovable-medium storage device, the term can also indirectly refer to the storage device associated with the volume. When multiple volumes are stored on a single storage medium transparently to the program, the volumes can be referred to as logical volumes.

volume group

A collection of either physical or logical volumes.

volume label

A unique identifier that a user assigns to a logical volume.

VPD See *vital product data*.

VSE/ESA

An IBM operating system, the letters of which represent virtual storage extended/enterprise systems architecture.

W

weight distribution area

The area that is required to distribute the weight of the storage unit.

worldwide node name (WWNN)

A unique 64-bit identifier for a host that contains a fibre-channel port. See also *worldwide port name*.

worldwide port name (WWPN)

A unique 64-bit identifier associated with a fibre-channel adapter port. It is assigned in an implementation- and protocol-independent manner. See also *worldwide node name*

write hit

A write operation in which the requested data is in the cache.

write penalty

The performance impact of a classical RAID-5 write operation.

WWNN

See *worldwide node name*.

WWPN

See *worldwide port name*.

X

xSeries

The product name of an IBM eServer product that emphasizes industry-standard server scalability and self-managing server technologies. It is the successor to the Netfinity family of servers.

Z

z/Architecture

An IBM architecture for mainframe computers and peripherals. The IBM eServer zSeries family of servers uses the z/Architecture architecture. It is the successor to the S/390 and 9672 family of servers. See also *iSeries*.

zoning

In fibre-channel environments, the grouping of multiple ports to form a virtual, private, storage network. Ports that are members of a zone can communicate with each other, but are isolated from ports in other zones.

z/OS An operating system for the IBM eServer product line that supports 64-bit real storage.

z/OS Global Mirror

A function of a storage server that assists a control program to maintain a consistent copy of a logical volume on another storage unit. All modifications of the primary logical volume by any attached host are presented in order to a single host. The host then makes these modifications on the secondary logical volume. This function was formerly called *extended remote copy* or *XRC*.

zSeries

An IBM eServer family of servers that emphasizes near-zero downtime.

IBM enterprise servers based on z/Architecture.

zSeries storage

Storage arrays and logical volumes that are defined in the DS8000 as connected to zSeries servers.

Index

A

- accessibility 615
- account setup 64
- add volumes to session, CLI 160
- adding ranks to extent pool, DS CLI 116
- address group 88
- AIX
 - CLI installation considerations 28
 - LIBPATH handling 28
- AIX LIBPATH 46
- alias volumes 445
- API, DS Open 12
- applykey 323
- arrays status 113
- arrays, creating (ckd) 101
- arrays, creating (FB) 86
- arrays, deleting 114
- assigning ranks 102
- audit report, console 595

B

- background copy 190

C

- calculating hexadecimal values 288
- changing volume states
 - Metro Mirror (DS CLI) 147, 156
- chckdvol 439
- checking command usage 58
- chextpool 419
- chfbvol 462
- chhostconnect 366
- chlcu 430
- chlss 458
- chpass 312
- chrank 408
- chsession 588
- chsi 349
- chsp 338
- chsu 343
- chuser 313
- chvolgrp 483
- Cisco MDS 9216 Multilayer Fabric Switch 27, 556, 574
- ckd
 - creating LCU 103
- CKD
 - deleting configurations 600
 - volume removal 447
- CKD configuration error 104
- CKD configuration error, correcting 106
- ckd volume
 - creating arrays 101
- CKD volume rank 102
- ckd volumes, creating 104
- clearvol 492
- CLI
 - calculating hexadecimal values 288

CLI (continued)

- displaying
 - WWNN 138
- pause 145
- postinstallation 54
- resuming Metro Mirror
 - relationship 144
- CLI commands on a command line 55
- CLI commands, overview 271
- CLI default configuration profile file 66
- CLI help 60
- CLI history function 58
- CLI interactive mode 58
- CLI, command-line interface 11
- command flags 271
- command line presentation 55
- commands
 - applykey 323
 - chckdvol 439
 - chextpool 419
 - chfbvol 462
 - chhostconnect 366
 - chlcu 430
 - chlss 458
 - chpass 312
 - chrank 408
 - chsession 588
 - chsi 349
 - chsp 338
 - chsu 343
 - chuser 313
 - chvolgrp 483
 - clearvol 492
 - commitflash 494
 - commitremoteflash 512
 - dscli 305
 - exit 307
 - failbackpprc 543
 - failoverpprc 546
 - freezepprc 562
 - help 308
 - lsaddressgrp 428
 - lsarray 400
 - lsarraysite 395
 - lsavailpprcport 529
 - lsckdvol 440
 - lsda 330
 - lsddm 334
 - lsxtpool 420
 - lsfbvol 463
 - lsflash 497
 - lsframe 329
 - lshba 332
 - lshostconnect 371
 - lshosttype 387
 - lshostvol 376
 - lsioport 357
 - lskey 324
 - lslcu 431
 - lslss 459
 - lsnetworkport 390

commands (continued)

- lsportprof 378
- lspprc 550
- lspprcpath 532
- lsrank 409
- lsremoteflash 517
- lsserver 350
- lssession 590
- lssi 352
- lststgencl 326
- lssu 344
- lsuser 315
- lsvolgrp 485
- managehostconnect 379
- managepwfile 316
- mkaliasvol 448
- mkarray 403
- mkckdvol 445
- mkesconpprcpath 535
- mkextpool 422
- mkfbvol 469
- mkflash 501
- mkgmir 574
- mkhostconnect 380
- mklcu 433
- mkpprc 556
- mkpprcpath 538
- mkrank 412
- mkremoteflash 521
- mksession 589
- mkuser 318
- mkvolgrp 486
- offloadauditlog 595
- overview structure 271
- pausegmir 575
- pausepprc 564
- quit 308
- resumegmir 577
- resumepprc 566
- resyncflash 495
- resyncremoteflash 514
- reverseflash 504
- revertflash 506
- revertremoteflash 523
- rarray 404
- rmckdvol 447
- rmextpool 423
- rmfbvol 472
- rmflash 507
- rmgmir 578
- rmhostconnect 383
- rmlcu 436
- rmpprc 568
- rmpprcpath 540
- rmrank 413
- rmremoteflash 525
- rmsession 594
- rmuser 320
- rmvolgrp 488
- setflashrevertible 510
- setioport 361

- commands (*continued*)
 - setnetworkport 388
 - setoutput 309
 - setremoteflashrevertible 527
 - setmpw 321
 - setvpn 340
 - showarray 404
 - showarraysite 397
 - showckdvol 449
 - showextpool 424
 - showfbvol 474
 - showgmir 580
 - showgmircg 585
 - showgmiroos 586
 - showhostconnect 384
 - showiport 362
 - showlcu 436
 - showlss 460
 - shownetworkport 393
 - showpass 321
 - showrank 414
 - shows 354
 - showsp 341
 - shows 346
 - showuser 322
 - showvolgrp 489
 - unfreezeflash 509
 - unfreezepprc 571
 - ver 325
 - commitflash 494
 - commitremoteflash 512
 - concurrent copy
 - modifying
 - timeout value (DS CLI) 148
 - conduit LSS 607
 - configuration
 - DS Storage Manager 10
 - reconfiguration 10
 - configuration error 92
 - Configuration overview 22
 - configuration profile file 66
 - configuring
 - fibre-channel I/O ports 95, 106
 - configuring I/O port, System i 95, 106
 - connection ID, caution 366
 - consistency group 185
 - inconsistencies 258
 - modifying
 - timeout value (DS CLI) 148
 - consistency group parameters, CLI 161
 - consistency groups
 - creating (DS CLI) 151
 - console audit report 595
 - console mode installation, DS CLI 37
 - controlling CLI page help 60
 - conversions
 - Global Copy
 - to Metro Mirror 610
 - Metro Mirror
 - to Global Copy 610
 - converting
 - Metro Mirror (DS CLI) 154
 - volume pairs 154
 - Copy Services
 - across 2105 and 1750 157
 - across 2105 and 2107 157
 - disaster recovery 19
 - Copy Services (*continued*)
 - licensed functions 21
 - overview 14
 - correcting
 - path-related configuration error 140
 - wrong value for remote WWNN 140
 - wrong value for target LSS ID 140
 - correcting a CKD configuration
 - error 106
 - correcting fixed block configuration
 - error 92
 - correcting FlashCopy relationships,
 - commitflash command 188
 - correcting FlashCopy relationships,
 - revertflash command 187
 - count key data volumes
 - creating extent pools 99
 - create FlashCopy relations, Global
 - Mirror 172
 - create Global Copy pairs 171
 - create paths, Global Mirror 169
 - creating
 - ckd volumes 104
 - consistency group (DS CLI) 143, 151
 - extent pools, fixed block 85
 - fixed block volume groups 92
 - fixed block volumes 88
 - FlashCopy relationship 125
 - LUN volumes 90
 - Metro Mirror volume pair
 - using a 2105 (DS CLI) 157
 - persistent FlashCopy 126
 - ranks for CKD volumes 102
 - SCSI host port connections 97
 - volume groups for System i 94
 - creating a rank, DS CLI 87
 - creating arrays (ckd) 101
 - creating arrays (FB) 86
 - creating extent pools, CKD 99
 - critical heavy mode 556
 - critical mode setting (DS CLI) 149
 - critical volume mode 556
- ## D
- DDM status 112
 - decimal 288
 - default CLI settings 66
 - default directories 55
 - default report format 309
 - delete arrays 114
 - delete extent pools, DS CLI 112
 - delete LCU 122
 - delete ranks 119
 - deleting
 - FlashCopy relationships 128
 - paths 141
 - single volume relationship (DS
 - CLI) 156
 - volume pair (DS CLI) 147
 - deleting CKD configurations 604
 - deleting fixed block configuration 600
 - delim report format 309
 - determining
 - available I/O ports 154
 - disaster recovery
 - Copy Services 19
 - disaster recovery (*continued*)
 - drain time 194
 - failback
 - scenario (DS CLI) 152
 - failover (DS CLI) 153
 - failover B volumes to A volumes 184
 - first pass 194
 - OOS 194
 - quiesce your system 194
 - restoring operations (planned) to the
 - intermediate site 236
 - restoring operations to the
 - intermediate site 212, 247
 - planned outage 205
 - restoring operations to the remote
 - site 217
 - restoring to the intermediate site 259
 - disaster recovery, Global Mirror
 - background copy 190
 - begin the failback recovery
 - process 192
 - checking status 183
 - consistency group analysis 185
 - correcting FlashCopy relationships,
 - commitflash command 188
 - end Global Mirror processing 182
 - fast reverse restore 189
 - Global Copy failback 197
 - Global Copy failover 197
 - process overview 181
 - reestablish FlashCopy relationships, B
 - to C 191
 - reestablishing paths 195
 - resume Global Mirror processing 198
 - resume Global Mirror processing at
 - the A site 198
 - resynchronize B volumes to A
 - volumes 193
 - revertflash and FlashCopy
 - relationships 187
 - disk drive module 112
 - Disk Manager
 - monitoring performance 12
 - displaying
 - paths 138
 - DS CLI
 - console mode installation 37
 - create paths 139
 - creating
 - volume pairs 142, 146
 - default directories 55
 - determining 154
 - first time login 55
 - graphical mode installation 31
 - I/O ports 154
 - installation on Linux 37
 - interactive mode 58
 - limitations 27
 - OpenVMS messages 613
 - persistent setup, OpenVMS 53
 - script mode 57
 - session limitation 27
 - silent mode installation 43
 - silent mode installation, DS CLI 43
 - single-shot mode 57
 - DS CLI command format 55

- DS CLI commands
 - System i 108
- DS CLI exit codes 61
- DS CLI installation
 - AIX considerations 28
 - default directories 28
 - HP Tru64 considerations 28
 - mounting the installation CD 28
 - preparing for 28
 - System i considerations 28
- DS CLI Tasks
 - Global Mirror
 - add volumes to session 160
 - create FlashCopy relationships 172
 - create Global Copy pairs 171
 - create session 173
 - end a session 175
 - end processing, no script 168
 - end processing, script mode 167
 - modify tuning parameters 161
 - pausing processing 165
 - query Global Mirror session 165
 - remove FlashCopy relationships 176
 - remove Global Copy relationships 177
 - remove paths 178
 - remove volumes 174
 - resume processing 166
 - setup, creating paths 169
 - start processing 167
 - view session 164
 - DS CLI Tasksmodify topology
 - Global Mirror 162
 - DS command-line interface 11
 - DS Open API 12
 - DS Storage Manager 10
 - DS8000 interfaces 10
 - DS8000 models 1
 - DS8100 series 1
 - DS8300 series 1
 - Model 922 1
 - Model 932 1
 - Model 9A2 1
 - Model 9B2 1
 - dscli 305
 - dscli command 305
 - during an unplanned outage 225

E

- end Global Mirror script mode 167
- ending Global Mirror processing 182
- ESCON 13
- exit 307
- exit codes, obtaining 61
- extent pool modifications, CLI 110
- extent pool names 85
- extent pool, adding ranks with DS
 - CLI 116
- extent pools, delete with DS CLI 112
- extent pools, System i 85

F

- failback
 - disaster recovery
 - scenario (DS CLI) 152
 - local to remote 608
 - restoring operations (planned) to the
 - intermediate site 236
 - restoring operations (unplanned) to
 - the intermediate site 247
 - using -force 236, 247
- failback recovery process 192
- failbackpprc 543
- failover
 - disaster recovery (DS CLI) 153
 - during an unplanned outage 225
 - planned outage
 - restoring operations to the
 - intermediate site 205
 - restoring operations (planned) to the
 - intermediate site 236
 - restoring operations (unplanned) to
 - the intermediate site 247
 - restoring operations to the
 - intermediate site 212, 259
 - restoring operations to the remote
 - site 217
 - using -force 236, 247
- failover and fallback 19
- failoverpprc 546
- fast reverse restore 189
- Fibre Channel Protocol 13
- fibre-channel I/O ports
 - configuring 95, 106
 - fc-al 95, 106
 - ficon 95, 106
 - scsi-fcp 95, 106
- FICON 13
- first pass 194
- first time login 55
- fixed block
 - deleting configurations 600
- fixed block configuration error,
 - correcting 92
- fixed block volume
 - creating arrays 86
- fixed block volume groups, creating 92
- fixed block volume groups,
 - modifying 599
- fixed block volumes
 - creating extent pools 85
- fixed block volumes, creating 88
- FlashCopy
 - background copy 133
 - creating a FlashCopy
 - relationship 125
 - deleting relationships 128
 - inconsistent
 - data on target volumes 258
 - persistent
 - creating 126
 - refreshing target volumes 130
 - reversing relationship 131
 - revertible option 132
 - revertible option defined 132
 - setflashrevertible 132
 - target volumes
 - on Metro Mirror source 135

- FlashCopy (*continued*)
 - target volumes (*continued*)
 - removing write inhibits 134
 - viewing information about FlashCopy
 - relationships 127
 - FlashCopy relationships, create 172
 - force parameter 543, 546
 - freezepprc 562
 - functions
 - Copy Services 14

G

- Global Copy
 - creating
 - volume pairs 146
- Global Copy failback 197
- Global Copy failover 197
- Global Copy pairs, create 171
- Global Mirror 19
- Global Mirror - CLI
 - add volumes to session 160
 - create FlashCopy relationships 172
 - create Global Copy pairs 171
 - create session 173
 - end processing, no script 168
 - end processing, script mode 167
 - end session 175
 - modify topology 162
 - modify tuning parameters 161
 - pausing processing 165
 - query session 165
 - remove FlashCopy relationships 176
 - remove Global Copy pair 177
 - remove paths 178
 - remove volumes 174
 - resume processing 166
 - setup environment, paths 169
 - start processing 167
 - view session 164
- Global Mirror processing, ending 182
- graphical mode installation, DS CLI 31

H

- help 308
- help, CLI 60
- hex 288
- hexadecimal 288
- host systems
 - adapter types 13
- HP-UX volume restriction 366

I

- I/O ports
 - available 154
- i5/OS
 - command format 57
 - green screen, DS CLI 108
 - library QDSCLI 108
 - removing DS CLI directly 75
 - silent mode CLI installation 43
- i5/OS password file directory 316
- IBMDSCLI\$STARTUP.COM
 - procedure 53

- implementation, RAID 9
- inband transactions 129, 607
- incremental FlashCopy
 - copying
 - updated tracks only 130
- incremental resynchronization
 - during a planned outage
 - at an intermediate site 259
- incrementalresync 556
- installing the CLI 54
- interactive mode
 - history function 58
 - setoutput command 58
- interfaces, DS8000 10
- iSeries
 - i5/OS password file directory 316
- issuing
 - DS CLI commands 108
 - i5/OS 108

J

- JVM not found error 46

K

- keyboards
 - accessibility features 615

L

- LCU creation 103
- LCU status 121
- LCU status, one 122
- LCU, modify with CLI 120
- LCUs, deleting 122
- legal
 - terms and conditions 618
- library QDSCLI 108
- licensed functions
 - Copy Services 21
- limitations
 - remote FlashCopy 129
- Linux
 - DS CLI installation, console mode 37
 - lshostvol command
 - considerations 376
- Linux password file directory 316
- lock account 313
- logical control unit creation 103
- logical control unit properties 122
- lsaddressgrp 428
- lsarray 400
- lsarraysite 395
- lsavailpprcport 529
- lsckdvol 440
- lsda 330
- lsddm 334
- lsxtpool 420
- lsfbvol 463
- lsflash 497
- lsframe 329
- lshba 332
- lshostconnect 371
- lshosttype 387

- lshostvol
 - Linux considerations 376
 - OpenVMS considerations 376
- lsioport 357
- lskey 324
- lslcu 431
- lslss 459
- lsnetworkport 390
- lsportprof 378
- lspprc 550
- lspprcpath 532
 - command 138
- lsrank 409
- lsremoteflash 517
- lsserver 350
- lssession 590
- lssi 352
 - command 138
- lsstgencl 326
- lssu 344
- lsuser 315
- lsvolgrp 485
- LUN sizes 469
- LUN volumes, creating 90

M

- managehostconnect 379
- management console
 - multiple units 8
 - overview 8
- managepwfile 316
- message types 613
- Metro Mirror
 - consistency groups 143
 - creating
 - volume pairs 142
 - FlashCopy target on Metro Mirror
 - source 135
 - scenario 608
- Metro/Global Mirror
 - restoring operations (planned) to the
 - intermediate site 236
 - restoring operations (unplanned) to
 - the intermediate site 247
 - setting up 201
- mkaliasvol 448
- mkarray 403
- mkckdvol 445
- mkconpprcpath 535
- mkextpool 422
- mkfbvol 469
- mkflash 501
- mkgmir 574
- mkhostconnect 380
- mklcu 433
- mkpprc 556
 - volume pairs 142, 146
- mkpprcpath 538
 - creating paths 139
- mkrank 412
- mkremoteflash 521
- mksession 589
- mkuser 318
- mkvolgrp 486
- Model 921 1, 4
- Model 922 5

- Model 92E 1
- Model 931 1, 4
- Model 932 5
- Model 9A2 5
- Model 9AE 1
- Model 9B2 1, 5
- modify
 - extent pool, CLI 110
- modify a rank, CLI 116
- modify an LCU, CLI 120
- modify topology, CLI 162
- modify tuning parameters, CLI 161
- modifying
 - concurrent copy timeout value (DS
 - CLI) 148
 - consistency group timeout (DS
 - CLI) 148
 - critical mode setting (DS CLI) 149
 - fixed block volume groups 599
 - timeout values (DS CLI) 148
 - z/OS Global Mirror (DS CLI) 149
- mounting the installation CD 28
- multiple LCU status 121
- multiple ranks status 118

N

- nonrevertible 187, 188
- Novell DS CLI graphical installation 31

O

- offloadauditlog 595
- OpenVMS
 - DS CLI message considerations 613
 - ECOs for installation 48
 - IBMDSCCLI\$STARTUP.COM
 - procedure 53
 - installing CLI 48
 - lshostvol command
 - considerations 376
 - mounting the installation CD 48
 - persistent setup 53
 - Polycenter software installation
 - utility 48
 - preparing for CLI installation 48
- operational limitations 27
- OS/400 DS CLI graphical installation 31
- os400all 94
- os400mask 94
- out-of-sync 194
- overview
 - Copy Services 14

P

- password file directories
 - i5/OS 316
 - LINUX 316
 - UNIX 316
 - Windows 316
- paths
 - creating 139
 - deleting 141
 - displaying 138
 - lspprcpath 151

- paths (*continued*)
 - mkpprcpath 139
 - port considerations 139
 - view existing 151
- Paths
 - correcting configuration error 140
- pausegmir 575
- pausepprc 564
 - suspending volumes 145
- performance gathering
 - Disk Manager 12
- persistent FlashCopy
 - creating 126
- persistent setup 53
- planned outage
 - restoring operations to the intermediate site 205
- Polycenter software installation
 - utility 48
- postinstallation
 - CLI 54
- postinstallation tasks 54
- primary volumes 184
- processing tips
 - 3 site recovery 556
 - cascade 556
 - incrementalresync 556
 - multiple volumes 556
 - volume grouping method 556

Q

- query Global Mirror session 165
- quiesce your system 194
- quit 308

R

- RAID
 - implementation 9
 - RAID 10 overview 10
 - RAID 5 overview 10
 - RAID overview 9
- rank configuration 87, 102
- rank group 1 or 1 102
- rank status, one 118
- rank, create with DS CLI 87
- rank, modify with CLI 116
- ranks status 118
- ranks, deleting 119
- reestablishing remote mirror and copy paths 195
- relationship
 - reversing FlashCopy 131
- release rank, CLI 116
- remote FlashCopy
 - limitations 129
- remote FlashCopy transactions 607
- Remote Mirror and Copy relationships
 - viewing (DS CLI) 153
- remove arrays 114
- remove LCUs 122
- remove ranks 119
- removing
 - uncorrupted incremental backup 134

- removing (*continued*)
 - write inhibits from FlashCopy target volumes 134
- removing DS CLI, i5/OS 75
- Replication
 - copy services functions 13
- report formats
 - controlling 58
 - default 309
 - delim 309
 - setoutput command 309
 - stanza 309
 - XML 309
- reserve rank, CLI 116
- restoring operations (planned) to the intermediate site 236
- restoring operations (unplanned) to the intermediate site 247
- restoring operations to the intermediate site 205, 212
- restoring operations to the remote site 217, 225
- resume Global Mirror processing 166
- resumegmir 577
- resumepprc 566
 - resuming volumes 144
- resync
 - target volumes in FlashCopy relationship 130
- resyncflash 495
- resynchronize B volumes to A volumes 193
- resyncremoteflash 514
- reverse a FlashCopy relationship 189
- reverseflash 504
- revertflash 506
- revertremoteflash 523
- rmarray 404
- rmckdvol 447
- rmextpool 423
- rmfbvol 472
- rmflash 507
- rmgmir 578
- rmhostconnect 383
- rmlcu 436
- rmpprc 568
- rmpprcpath 540
- rmrank 413
- rmremoteflash 525
- rmsession 594
- rmuser 320
- rmvolgrp 488

S

- script mode 57
- script mode, end processing 167
- SCSI host port connections
 - creating 97
- serial number, obtaining 59
- setflashrevertible 510
- setioport 361
- setnetworkport 388
- setoutput 309
- setremoteflashrevertible 527
- setmpw 321

- setting up
 - Metro/Global Mirror 201
- setup, account 64
- setvprn 340
- showarray 404
- showarraysite 397
- showckdvol 449
- showextpool 424
- showfbvol 474
- showgmir 580
- showgmircg 585
- showgmiros 586
- showhostconnect 384
- showioport 362
- showlcu 436
- showlss 460
- shownetworkport 393
- showpass 321
- showrank 414
- showsi 354
- showsp 341
- showsu 346
- showuser 322
- showvolgrp 489
- single volume relationship
 - deleting (DS CLI) 156
- single-shot mode
 - command format, i5/OS 57
 - command format, Windows 57
- stanza report format 309
- starting
 - FlashCopy 133
- status information, fixed block
 - volume 463
- stop the process
 - Global Mirror 578
- storage image ID, obtaining 59
- System i
 - bypassing uninstaller 75
 - CLI installation considerations 28
 - configuring I/O port 95, 106
 - creating
 - FlashCopy relationships 125
 - Global Copy pairs 146
 - Metro Mirror pairs 142
 - DS CLI commands 108
 - extent pool considerations 85
 - host connections 97
 - interactive mode 58
 - issuing 108
 - LUN volumes 90
 - lunmapping 489
 - os400all 94
 - os400mask 94
 - postinstallation tasks 54
 - removing CLI, direct method 75
 - removing CLI, remote method 75
 - script mode 57
 - serial number reporting 354
 - showvolgrp 489
 - single-shot mode 57
 - volume group creation
 - external load source consideration 94
 - multipath consideration 94

T

- target volume
 - commit changes 137
 - discard changes 136
 - FlashCopy
 - commit changes 137
 - discard changes 136
- thaw operation 205, 212, 217, 225, 236, 247, 571, 608
- timeout connection 66
- timeout values
 - modifying
 - concurrent copy (DS CLI) 148
 - consistency group (DS CLI) 148
 - critical mode enabled (DS CLI) 148
 - z/OS Global Mirror (DS CLI) 148
- topology defined, CLI 162
- TotalStorage Productivity Center
 - copy services 13
 - for Disk 12
 - Replication 13
- Trademarks 619
- tuning parameters, CLI 161

U

- unassign rank, CLI 116
- unattended (silent) mode installation 43
- unattended (silent) mode installation, DS CLI 43
- unfreezeflash 509
- unfreezepprc 571
- uninstall.dat 75
- uninstall.jar 75
- UNIX DS CLI graphical installation 31
- UNIX password file directory 316
- unlock account 313

V

- ver 325
- view a specific LCU 122
- view a specific rank 118
- view array sites status 113
- view arrays status 113
- view LCU status 121
- view ranks status 118
- view session, CLI 164
- viewing
 - existing paths 151
 - information about FlashCopy relationships 127
 - Remote Mirror and Copy relationships (DS CLI) 153
- volume format processing options 556
- volume pair
 - deleting (DS CLI) 147
- volume removal, CKD 447
- volumes
 - FlashCopy target volumes
 - commit changes 137
 - discard changes 136
- volumes, alias 445

W

- Windows DS CLI graphical installation 31
- Windows password file directory 316
- write acceleration restriction 27, 556, 574
- write inhibits
 - mkflash command 134
 - removing from FlashCopy target volumes 134
- wwnn
 - displaying for paths 138

X

- XML report format 309

Z

- z/OS Global Mirror
 - modifying
 - critical mode setting (DS CLI) 149
 - timeout value (DS CLI) 149
- zones
 - WWPN 13

Readers' comments — we would like to hear from you

IBM System Storage DS8000
Command-Line Interface
User's Guide
Version 1 Release 2

Publication No. SC26-7916-01

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Submit your comments using one of these channels:

- Send your comments to the address on the reverse side of this form.

If you would like a response from IBM, please fill in the following information:

Name

Address

Company or Organization

Phone No.

E-mail address



Fold and Tape

Please do not staple

Fold and Tape



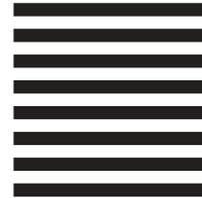
NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department 61C
9032 South Rita Road
Tucson, AZ 85775-4401



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

SC26-7916-01

