



IBM System Storage SAN Volume Controller and Storwize V7000 Command-Line Interface User's Guide

Version 6.1.0

GC27-2287-00





IBM System Storage SAN Volume Controller and Storwize V7000 Command-Line Interface User's Guide

Version 6.1.0

GC27-2287-00

Note

Before using this information and the product it supports, read the information in “Notices” on page 503.

This edition applies to the IBM System Storage SAN Volume Controller System Storage SAN Volume Controller and Storwize V7000, Version 6.1.0, and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces SC26-7903-06.

© **Copyright IBM Corporation 2003, 2010.**

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

Contents

Tables ix

About this guide. xi

Who should use this guide	xi
Summary of changes	xi
Summary of changes for GC27-2287-00 SAN Volume Controller Command-Line Interface User's Guide	xi
Summary of changes for SC26-7903-05 and SC26-7903-06 SAN Volume Controller Command-Line Interface User's Guide	xv
Emphasis	xi
SAN Volume Controller library and related publications	xix
How to order IBM publications	xxi
Sending your comments	xxi
Syntax diagrams	xxi
Terminology	xxiii
CLI special characters	xxiii
Using wildcards in the SAN Volume Controller CLI	xxiii
Data types and value ranges	xxiv
CLI commands and parameters	xxix
CLI flags	xxx
CLI messages	xxxi

Chapter 1. Preparing the SSH client. 1

Preparing the SSH client on a Windows host	1
Generating an SSH key pair using PuTTY	2
Configuring a PuTTY session for the CLI	3
Connecting to the CLI using PuTTY	3
Starting a PuTTY session for the CLI	5
Preparing the SSH client on an AIX or Linux host	6
Generating an SSH key pair using OpenSSH.	6
Connecting to the CLI using OpenSSH	7
Creating users	7

Chapter 2. Copying the SAN Volume Controller software upgrade files using PuTTY scp 9

Chapter 3. Using the CLI 11

Setting the cluster time using the CLI.	11
Viewing and updating license settings using the CLI	12
Displaying cluster properties using the CLI.	12
Maintaining passwords for the front panel using the CLI	13
Re-adding a repaired node to a cluster using the CLI	14
Displaying node properties using the CLI	17
Discovering MDisks using the CLI	18
Creating MDisk groups (storage pools) using the CLI	19

Adding MDisks to MDisk groups (storage pools) using the CLI.	21
Setting a quorum disk using the CLI	22
Modifying the amount of available memory for Copy Services and VDisk Mirroring features using the CLI.	23
Creating VDisks (volumes) using the CLI	24
Adding a copy to a VDisk (volume) using the CLI	27
Deleting a copy from a VDisk (volume) using the CLI	28
Configuring host objects using the CLI	28
Creating VDisk-to-host mappings using the CLI	30
Creating FlashCopy mappings using the CLI	30
Preparing and starting a FlashCopy mapping using the CLI.	31
Stopping FlashCopy mappings using the CLI	32
Deleting a FlashCopy mapping using the CLI	32
Creating a FlashCopy consistency group and adding mappings using the CLI	32
Preparing and starting a FlashCopy consistency group using the CLI	34
Stopping a FlashCopy consistency group using the CLI.	35
Deleting a FlashCopy consistency group using the CLI.	35
Creating Metro Mirror or Global Mirror relationships using the CLI	36
Modifying Metro Mirror or Global Mirror relationships using the CLI	36
Starting and stopping Metro Mirror or Global Mirror relationships using the CLI.	36
Displaying the progress of Metro Mirror or Global Mirror relationships using the CLI	37
Switching Metro Mirror or Global Mirror relationships using the CLI	37
Deleting Metro Mirror and Global Mirror relationships using the CLI	37
Creating Metro Mirror or Global Mirror consistency groups using the CLI	38
Modifying Metro Mirror or Global Mirror consistency groups using the CLI	38
Starting and stopping Metro Mirror or Global Mirror consistency-group copy processes using the CLI.	38
Deleting Metro Mirror or Global Mirror consistency groups using the CLI	39
Creating Metro Mirror and Global Mirror partnerships using the CLI	39
Modifying Metro Mirror and Global Mirror partnerships using the CLI	39
Starting and stopping Metro Mirror and Global Mirror partnerships using the CLI.	40
Deleting Metro Mirror and Global Mirror partnerships using the CLI	40
Determining the WWPNS of a node using the CLI	40

Listing node-dependent VDisks (volumes) using the CLI	41
Determining the VDisk name from the device identifier on the host	41
Determining the host that a VDisk (volume) is mapped to.	42
Determining the relationship between VDisks (volumes) and MDisks using the CLI.	42
Determining the relationship between MDisks and controller LUNs using the CLI	43
Increasing the size of your cluster using the CLI	43
Adding a node to increase the size of a cluster using the CLI.	43
Migrating a VDisk to a new I/O group using the CLI	44
Validating and repairing mirrored VDisk copies using the CLI.	45
Repairing a space-efficient VDisk using the CLI	47
Recovering from offline VDisks using the CLI	47
Recovering a node and returning it to the cluster using the CLI.	48
Recovering offline VDisks using the CLI.	49
Moving offline VDisks to their original I/O group using the CLI	50
Recording WWPN changes of replaced host HBAs	50
Expanding VDisks (volumes) using the CLI	51
Expanding a VDisk (volume) that is mapped to an AIX host	52
Expanding a VDisk that is mapped to a Microsoft Windows host using the CLI	52
Shrinking a virtual disk using the CLI	53
Migrating extents using the CLI	53
Migrating VDisks (volumes) between MDisk groups using the CLI.	55
Migrating a VDisk between I/O groups using the CLI	56
Creating an image mode VDisk (volume) using the CLI	57
Migrating data to an image mode virtual disk using the CLI.	57
Deleting a node from a cluster using the CLI	58
Performing the cluster maintenance procedure using the CLI.	60
Modifying the cluster IP addresses using the CLI.	61
Changing the cluster gateway address using the CLI	62
Changing the relationship bandwidth for a cluster using the CLI.	62
Configuring the cluster for iSCSI using the CLI	62
Configuring or modifying an iSCSI alias using the CLI.	64
Configuring the iSNS server address using the CLI	64
Configuring cluster iSCSI authentication using the CLI.	64
Configuring remote authentication service using CLI	65
Creating and working with user groups using the CLI	66
Creating and working with users using the CLI	67
Setting up SNMP notifications using the CLI	68
Setting up syslog notifications using the CLI	69

Setting up e-mail event notifications and inventory reports using the CLI	70
Setting up e-mail servers using the CLI	71
Changing cluster passwords using the CLI	71
Changing the locale setting using the CLI	72
Viewing the feature log using the CLI	72
Analyzing the error log using the CLI	72
Shutting down a cluster using the CLI	72
Upgrading the software automatically using the CLI	73

Chapter 4. Audit log commands 77

catauditlog	77
dumppauditlog	78
lsauditlogdumps (Deprecated)	79

Chapter 5. User management commands 81

chauthservice.	81
chcurrentuser.	83
chuser	84
chusergrp	85
mkuser	86
mkusergrp.	87
rmuser	88
rmusergrp.	88

Chapter 6. Email and event notification commands 91

chemail.	92
chemailserver.	93
chemailuser	94
chsnmpserver	95
chsyslogserver	96
mkemailserver	97
mkemailuser	98
mksnmpserver	99
mksyslogserver.	101
rmemailserver	102
rmemailuser.	102
rmsnmpserver	103
rmsyslogserver	103
sendinventoryemail	104
startemail.	104
stopemail.	105
testemail	105

Chapter 7. Cluster commands 107

addnode	108
cfgportip	110
chcluster	111
chclusterip	113
chiogrp	115
chnode	118
cleardumps	119
cpdumps	120
detectmdisk	121
ping	123
rmnode	123
rmportip	125
setclustertime	126

setpwdreset	127
settimezone	127
startstats	128
stopcluster	130
stopstats (Deprecated)	132

Chapter 8. Backup and restore commands 133

backup	133
clear	133
help	134
restore.	135

Chapter 9. Cluster diagnostic and service-aid commands 137

applysoftware	137
cherrstate.	138
clearerrlog	139
dumperrlog	139
finderr	140
setlocale	140
svqueryclock	141
writesernum.	142

Chapter 10. Host commands 143

addhostiogr	143
addhostport	143
chhost	144
mkhost	146
rmhost	147
rmhostiogr	148
rmhostport	149

Chapter 11. Virtual disk commands 151

addvdiskcopy	151
chvdisk	155
expandvdisksize	158
mkvdisk	160
mkvdiskhostmap	166
recovervdisk.	168
recovervdiskbycluster	168
recovervdiskbyiogr	169
repairsevdiskcopy	170
repairvdiskcopy	171
rmvdisk	172
rmvdiskcopy	173
rmvdiskhostmap	174
shrinkvdisksize.	175
splitvdiskcopy	177

Chapter 12. Managed disk group commands 179

addmdisk	179
chmdiskgrp	180
mkmdiskgrp	181
rmmdisk	183
rmmdiskgrp.	184

Chapter 13. Managed disk commands 187

applymdisksoftware (Discontinued)	187
chmdisk	187
chquorum	187
dumpallmdiskbadblocks.	189
dumpmdiskbadblocks	190
includemdisk	191
setquorum (Deprecated).	192
triggermdiskdump (Discontinued)	192

Chapter 14. IBM FlashCopy commands 193

chfcconsistgrp	193
chfcmap	193
mkfcconsistgrp	195
mkfcmap	196
prestartfcconsistgrp	198
prestartfcmap	199
rmfcconsistgrp	200
rmfcmap	200
startfcconsistgrp	201
startfcmap	202
stopfcconsistgrp	203
stopfcmap	204

Chapter 15. Metro Mirror and Global Mirror commands 207

chpartnership	207
chrconsistgrp	208
chrcrelationship	208
mkpartnership	210
mkrconsistgrp	211
mkrcrelationship	212
rmpartnership	214
rmrconsistgrp	214
rmrcrelationship	215
startrrconsistgrp	216
startrrcrelationship	218
stoprrconsistgrp	220
stoprrcrelationship	221
switchrconsistgrp.	222
switchrcrelationship	223

Chapter 16. Migration commands. . . 225

migrateexts	225
migratetoimage.	226
migratevdisk	227

Chapter 17. Tracing commands. . . . 229

setdisktrace	229
settrace	229
starttrace	231
stoptrace	232

Chapter 18. Drive commands 233

applydrivesoftware	233
chdrive	233
lsdrive.	234
lsdrivelba.	236
lsdriveprogress.	237

triggerdrivedump	238
Chapter 19. Enclosure commands	239
chenclosure	239
chenclosurecanister	239
chenclosureslot	240
lsenclousure	241
lsenclousurebattery	243
lsenclousurecanister	244
lsenclousurepsu	247
lsenclousureslot	248
triggerenclosuredump	250

Chapter 20. Attributes of the -filtervalue parameters 251

Chapter 21. Overview of the dumps commands 261

Chapter 22. Information commands 263

caterrlog (Deprecated)	263
caterrlogbyseqnum (Deprecated)	263
ls2145dumps (Deprecated)	263
lscimomdumps (Deprecated)	263
lscopystatus	263
lscluster	264
lsclustercandidate	268
lsclusterrip	269
lscontroller	271
lscontrollerdependentvdisk	273
lscurrentuser	274
lsdiscoverystatus	275
lsdumps	275
lsemailer	277
lsemaileruser	278
lserrlogbyfcconsistgrp (Deprecated)	279
lserrlogbyfcmap (Deprecated)	279
lserrlogbyhost (Deprecated)	279
lserrlogbyiogrp (Deprecated)	279
lserrlogbymdisk (Deprecated)	279
lserrlogbymdiskgrp (Deprecated)	279
lserrlogbynode (Deprecated)	279
lserrlogbyrconsistgrp (Deprecated)	279
lserrlogbyrrelationship (Deprecated)	280
lserrlogbyvdisk (Deprecated)	280
lserrlogdumps (Deprecated)	280
lseventlog	280
lsfabric	284
lsfcconsistgrp	285
lsfcmap	287
lsfcmapcandidate	289
lsfcmapprogress	290
lsfcmapdependentmaps	291
lsfeaturedumps (Deprecated)	292
lsfreeextents	292
lshbaportcandidate	293
lshost	294
lshostiogrp	296
lshostvdiskmap	297
lsiogrp	298

lsiogrp	300
lsiogrpcandidate	301
lsiostatsdumps (Deprecated)	302
lsiotracedumps (Deprecated)	302
lssiscsiauth	302
lslicense	303
lsmdisk	304
lsmdiskdumps (Deprecated)	309
lsmdisklba	309
lsmdiskcandidate	310
lsmdiskextent	311
lsmdiskgrp	313
lsmdiskmember	316
lsmigrate	318
lsnode	319
lsnodecandidate	323
lsnodeindependentvdisk (Deprecated)	324
lsnodevpd	324
lsportip	328
lsquorum	330
lsrconsistgrp	331
lsrrelationship	334
lsrrelationshipcandidate	336
lsrrelationshipprogress	337
lsrepairsevdiskcopyprogress	338
lsrepairvdiskcopyprogress	339
lsrmvdiskdependentmaps	341
lsroute	342
lssevdiskcopy	342
lssnmpserver	346
lssoftwaredumps (Deprecated)	347
lssoftwareupgradestatus	347
lssyslogserver	347
lstimezones	349
lsuser	349
lsusergrp	351
lsvdisk	352
lsvdiskcopy	357
lsvdiskdependentmaps	361
lsvdiskextent	361
lsvdiskfcmapcopies	363
lsvdiskfc mappings	363
lsvdiskhostmap	364
lsvdisklba	365
lsvdiskmember	366
lsvdiskprogress	368
lsvdisksyncprogress	369
lsdependentvdisk	370
lssasfabric	371
showtimezone	372
resetleds	373

Chapter 23. Licensing commands 375

chlicense	375
dumpinternallog	376

Chapter 24. Service mode commands (Discontinued) 379

applysoftware (Discontinued)	379
cleardumps (Discontinued)	379

dumperrlog (Discontinued)	379
exit (Discontinued)	379

Chapter 25. Service mode information commands (Discontinued) 381

ls2145dumps (Discontinued)	381
lscimomdumps (Discontinued)	381
lsclustervpd (Discontinued)	381
lserrlogdumps (Discontinued)	381
lsfeaturedumps (Discontinued)	381
lsiostatsdumps (Discontinued)	381
lsiotracedumps (Discontinued)	381
lsmdiskdumps (Discontinued)	381
lsnodevpd (Discontinued)	382
lssoftwaredumps (Discontinued)	382

Chapter 26. Array commands 383

charray	383
charraymember	383
lsarray	385
lsarrayinitprogress	388
lsarraylba	389
lsarraymember	390
lsarraymembergoals	392
lsarraymemberprogress	393
lsarraysyncprogress	394
mkarray	395
recoverarray	397
recoverarraybycluster	397
rmarrray	398

Chapter 27. Service information commands 399

chnodeled	399
lscmdstatus	399

lsfiles	400
lsservicenodes	400
lsservicerecommendation	402
lsservicestatus	402

Chapter 28. Service task commands 407

chenclosurevpd	407
chserviceip	408
chwwnn	409
cpfiles	410
installsoftware	411
leavecluster	411
metadata	412
mkcluster	413
rescueneode	414
resetpassword	414
snap	414
setpacedccu	415
settempsshkey	415
startservice	416
stopnode	416
stopservice	417
t3recovery	417

Chapter 29. Controller command . . . 419

chcontroller	419
------------------------	-----

Appendix. Accessibility 501

Notices 503

Trademarks	505
----------------------	-----

Index 507

Tables

1.	Terminology mapping table	xi	24.	lsenclosurecanister output	245
2.	Other IBM publications	xx	25.	lsenclosurepsu output.	247
3.	IBM documentation and related websites	xx	26.	lsenclosureslot output.	249
4.	Maximum VDisk (volume) capacity by extent size	20	27.	Valid filter attributes	251
5.	Memory required for VDisk Mirroring and Copy Services.	23	28.	Attribute values.	265
6.	RAID level comparisons	24	29.	lseventlog output	281
7.	VDisk (volume) copy resynchronization rates	25	30.	MDisk output	306
8.	ip_address_list formats	115	31.	lsmdisklba command output	310
9.	Memory required for VDisk Mirroring and Copy Services	117	32.	svcinfo lsnode attribute values.	320
10.	RAID level comparisons	117	33.	svcinfo lsnode attribute values.	321
11.	Storage pool Easy Tier settings.	153	34.	lsnodecandidate outputs	324
12.	Relationship between the <i>rate</i> value and the data copied per second	154	35.	lsnodecandidate outputs	324
13.	Storage pool Easy Tier settings.	163	36.	lsrcconsistgrp command output values	333
14.	Relationship between the <i>rate</i> value and the data copied per second	165	37.	lsrcrelationship command attributes and values	335
15.	Number of extents reserved by extent size	188	38.	lsvdisklba command output scenarios	366
16.	Relationship between the <i>rate</i> , data rate and grains per second values.	195	39.	lssasfabric output	371
17.	Relationship between the <i>rate</i> , data rate and grains per second values.	198	40.	charraymember combination options	384
18.	stoprcconsistgrp consistency group states	220	41.	MDisk output	387
19.	stoprcrelationship consistency group states	222	42.	lsarrayinitprogress output	389
20.	lsdrive output	235	43.	lsarraylba output	390
21.	lsdrivelba output	237	44.	lsarraymemberoutput	391
22.	lsenclosure output	242	45.	lsarraymembergoals output	392
23.	lsenclosurebattery outputs	244	46.	lsarraymemberprogress output.	394
			47.	lsarraysyncprogress output	395
			48.	lsservicenodes outputs	401
			49.	lsservicenodes outputs	401
			50.	lsservicestatus output	403
			51.	lsservicestatus output	404

About this guide

This publication provides information that helps you configure and use the IBM® System Storage® SAN Volume Controller SAN Volume Controller and Storwize® V7000.

Who should use this guide

This guide is intended for system administrators or others who install and use the SAN Volume Controller or Storwize V7000.

Before you use the SAN Volume Controller, you should have an understanding of storage area networks (SANs), the storage requirements of your enterprise, and the capabilities of your storage units.

Summary of changes

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. This summary of changes describes new functions that have been added to this release.

Summary of changes for GC27-2287-00 SAN Volume Controller Command-Line Interface User's Guide

This topic describes the changes that have been made to the SAN Volume Controller *Command-Line Interface User's Guide* since the previous edition (SC26-7903-06).

6.1.0 Terminology changes

To coincide with new and existing IBM products and functions, several common terms have changed and are incorporated in the SAN Volume Controller information. Certain SAN Volume Controller information, particularly command-line interface (CLI) documentation, remains primarily unchanged.

The following table shows the current and previous usage of the changed common terms.

Table 1. Terminology mapping table

6.1.0 SAN Volume Controller term	Previous SAN Volume Controller term	Description
event	error	An occurrence of significance to a task or system. Events can include completion or failure of an operation, a user action, or the change in state of a process.
host mapping	VDisk-to-host mapping	The process of controlling which hosts have access to specific volumes within a cluster.
storage pool	managed disk (MDisk) group	A collection of storage capacity that provides the capacity requirements for a volume.
thin provisioning (or thin-provisioned)	space-efficient	The ability to define a storage unit (full system, storage pool, volume) with a logical capacity size that is larger than the physical capacity assigned to that storage unit.
volume	virtual disk (VDisk)	A discrete unit of storage on disk, tape, or other data recording medium that supports some form of identifier and parameter list, such as a volume label or input/output control.

| **New information**

The following new commands have been added for this edition:

- applydrivesoftware
- charray
- charraymember
- chdrive
- chemail
- chenclosure
- chenclosurecanister
- chenclosureslot
- chenclosurevpd
- chnodeled
- chserviceip
- chwwnn
- cpfiles
- dumperrlog
- dumpallmdiskbadblocks
- dumpmdiskbadblocks
- installsoftware
- leavecluster
- lsarray
- lsarraylba
- lsarraymember
- lsarraymembergoals
- lsarraymemberprogress
- lsarraysyncprogress
- lscmdstatus
- lsdependentvdisk
- lsdiscoverystatus
- lsdrive
- lsdrivelba
- lsdriveprogress
- lsdumps
- lsenclature
- lseventlog
- lsenclaturebattery
- lsenclaturecanister
- lsenclaturepsu
- lsenclatureslot
- lsfiles
- lssasfabric
- lsservicenodes
- lsservicestatus
- lsservicerecommendation

- metadata
- mkarray
- mkcluster
- recovervdisk
- recoverarray
- recoverarraybycluster
- rescuenode
- resetleds
- resetpassword
- rrray
- settempsshkey
- setlocale
- setpacedccu
- snap
- startservice
- stopnode
- stopservice
- t3recovery
- triggerdrivedump
- triggerenclosedump

Changed information

The following commands and topics have been updated for this edition:

- addmdisktomdiskgroup
- addnode
- addnodetocluster
- addvdiskcopy
- applydrivesoftware
- applysoftwareupgrade
- backup
- catauditlog
- changelocale
- chenclosureslot
- chcluster
- chclusterip
- chlicense
- chvdisk
- chiogrp
- chnode
- createmdiskgroup
- createsquentialorstripedmodevdisk
- Data types and value ranges
- discoverystatus
- drivelba
- ls2145dumps

- lsarray
- lsarraylba
- lsarraymember
- lsauditlog
- lscluster
- lsclusterip
- lscontroller
- lsdiscoverystatus
- lsdrive
- lsdrivelba
- lsenclurepsu
- lsfabric
- lsiogrp
- lsio tracedumps
- lslicense
- lsmdisk
- lsmdiskgrp
- lsnode
- lsnodecandidate
- lsnodevpd
- lsquorum
- lssevdiskcopy
- lsvdisk
- lsvdiskcopy
- lsvdisklba
- migratetoimage
- mkemailserver
- mkmdiskgrp
- mkvdisk
- modifycluster
- modifyiogroup
- modifymdisk
- modifymdiskgroup
- modifyvdisk
- recoverydisk
- restore
- setlocale
- startstats
- stopstats
- Viewing and updating license settings using the CLI

Deprecated and discontinued information

The following commands have been deprecated or discontinued:

- addnodetocluster
- applymdisksoftware

- caterrlog
- caterrlogbyseqnum
- chquorum
- ls2145dumps
- lscimomdumps
- lscontrollerdependentvdisks
- lsnodedependentvdisks
- lserrlogbyfcconsistgrp
- lserrlogbyrcrelationship
- lserrlogbyfcmap
- lserrlogbymdisk
- lserrlogbymdiskgroup
- lserrlogbyhost
- lserrlogbyiogrp
- lserrlogbynode
- lserrlogbyrconsistgrp
- lserrlogbyvdisk
- lserrlogdumps
- lsfeaturedumps
- lsiostatsdumps
- lsio tracedumps
- lsmdiskdumps
- lsnodedependentvdisks
- lssoftwaredumps
- setquorum
- stopstats
- triggermdiskdump
- Service mode commands
- Service mode information commands

Summary of changes for SC26-7903-05 and SC26-7903-06 SAN Volume Controller Command-Line Interface User's Guide

This topic describes the changes that have been made to the SAN Volume Controller *Command-Line Interface User's Guide* since the previous edition (SC26-7903-04).

New information

The following new SAN Volume Controller commands have been added for this edition:

- Cluster commands:
 - cfgportip
 - chclusterip
 - ping
 - rmpportip
- Cluster diagnostic and service-aid commands:
 - svctask applysoftware
 - svctask setlocale

- svctask writesernum
- E-mail and event notification commands:
 - chemailserver
 - chsnmpserver
 - chsyslogserver
 - mkemailserver
 - mksnmpserver
 - mksyslogserver
 - rmemailserver
 - rmsnmpserver
 - rmsyslogserver
- Information commands:
 - lsclusterip
 - lscurrentuser
 - lsemailserver
 - lsiscsiauth
 - lsmdiskdumps
 - lsnodedependentvdisks
 - lspportip
 - lsquorum
 - lsrmdiskdependentmaps
 - lsroute
 - lssnmpserver
 - lssyslogserver
 - lsuser
 - lsusergrp
 - lsvdiskfmapcopies
- Managed disk commands:
 - applymdisksoftware
 - triggermdiskdump
- Service mode information commands:
 - lsmdiskdumps
- User management commands:
 - chauthservice
 - chcurrentuser
 - chuser
 - chusergrp
 - mkuser
 - mkusergrp
 - rmuser
 - rmusergrp

Changed information

The following commands and topics have been updated for this edition:

- Cluster commands:

- chcluster
- chnode
- cleardumps
- cpdumps
- rmnode
- setpwdreset
- startstats
- stopcluster
- Cluster diagnostic and service-aid commands:
 - applysoftware
 - setlocale
 - setevent
 - writesernum
- E-mail and event notification commands:
 - chemail
 - chemailuser
 - mkemailuser
 - rmemailuser
- FlashCopy[®] commands:
 - mkfcmap
 - prestartfcconsistgrp
 - prestartfcmap
 - rmfcmap
 - startfcconsistgrp
 - startfcmap
 - stopfcconsistgrp
 - stopfcmapping
- Host commands:
 - addhostport
 - chhost
 - mkhost
 - rmhostiogr
 - rmhostport
- Information commands:
 - lscluster
 - lsclustercandidate
 - lsclusterip
 - lscontroller
 - lsemailuser
 - lsfcmap
 - lshost
 - lshostvdiskmap
 - lsioatdumps
 - lsmdisk
 - lsmdiskdumps

- lsmdiskgrp
- lsmdisklba
- lsnode
- lsnodecandidate
- lsnodevpd
- lssoftwareupgradestatus
- lsvdiskhostmap
- lsvdisksyncprogress
- Licensing commands:
 - chlicense
- Managed disk commands:
 - setquorum
- Managed disk group commands:
 - addmdisk
 - mkmdiskgrp
- Metro Mirror and Global Mirror commands:
 - mkpartnership
 - rmpartnership
 - startrcconsistgrp
 - startrcrelationship
- Service mode commands:
 - applysoftware
- Service mode Information commands:
 - lsclustervpd
 - lsnodevpd
 - lsioatdumps
 - lsmdiskdumps
- Virtual disk commands:
 - chvdisk
 - lsvdisklba
- Topics:
 - Attributes of the filtervalue parameters
 - CLI commands and parameters
 - Configuring a PuTTY session for the CLI
 - Error log commands
 - Secure Shell
 - Generating an SSH key pair using PuTTY

Deprecated information

The following SAN Volume Controller commands have been discontinued:

- Cluster diagnostic and service-aid commands:
 - svcservicetask addnode
 - svcservicetask applysoftware
 - svcservicetask rmnode
 - svcservicetask setlocale

- svcservicetask writesernum
- E-mail and event notification commands:
 - setemail
- Error log commands:
 - svcservicetask cherrstate
 - svcservicetask clearerrlog
 - svcservicetask dumperrlog
 - svcservicetask finderr
 - svcservicetask setevent
- Information commands:
 - lsnodes
 - lsshkeys
- Role-based security commands:
 - mkauth
 - rmath
 - lsauth
- Secure Shell key commands:
 - addsshkey
 - rmallsshkeys
 - rmsshkey

Emphasis

Different typefaces are used in this guide to show emphasis.

The following typefaces are used to show emphasis:

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

SAN Volume Controller library and related publications

Product manuals, other publications, and websites contain information that relates to SAN Volume Controller.

SAN Volume Controller Information Center

The IBM System Storage SAN Volume Controller Information Center contains all of the information that is required to install, configure, and manage the SAN Volume Controller. The information center is updated between SAN Volume Controller product releases to provide the most current documentation. The information center is available at the following website:

publib.boulder.ibm.com/infocenter/svc/ic/index.jsp

SAN Volume Controller library

Unless otherwise noted, the publications in the SAN Volume Controller library are available in Adobe portable document format (PDF) from the following website:

Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

Other IBM publications

Table 2 lists IBM publications that contain information related to the SAN Volume Controller.

Table 2. Other IBM publications

Title	Description	Order number
<i>IBM System Storage Productivity Center Introduction and Planning Guide</i>	This guide introduces the IBM System Storage Productivity Center hardware and software.	SC23-8824
<i>Read This First: Installing the IBM System Storage Productivity Center</i>	This guide describes how to install the IBM System Storage Productivity Center hardware.	GI11-8938
<i>IBM System Storage Productivity Center User's Guide</i>	This guide describes how to configure the IBM System Storage Productivity Center software.	SC27-2336
<i>IBM System Storage Multipath Subsystem Device Driver User's Guide</i>	This guide describes the IBM System Storage Multipath Subsystem Device Driver for IBM System Storage products and how to use it with the SAN Volume Controller.	GC52-1309

IBM documentation and related websites

Table 3 lists websites that provide publications and other information about the SAN Volume Controller or related products or technologies.

Table 3. IBM documentation and related websites

Website	Address
Support for SAN Volume Controller (2145)	Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145
Support for IBM System Storage and IBM TotalStorage products	www.ibm.com/storage/support/
IBM Publications Center	www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss
IBM Redbooks® publications	www.redbooks.ibm.com/

Related accessibility information

To view a PDF file, you need Adobe Acrobat Reader, which can be downloaded from the Adobe website:

www.adobe.com/support/downloads/main.html

How to order IBM publications

The IBM 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 at no 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 website:

www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss

Sending your comments

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

To submit any comments about this book or any other SAN Volume Controller documentation:

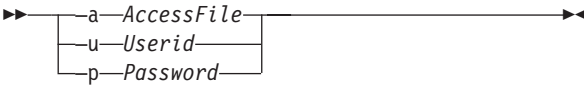


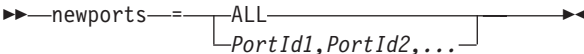

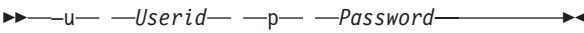
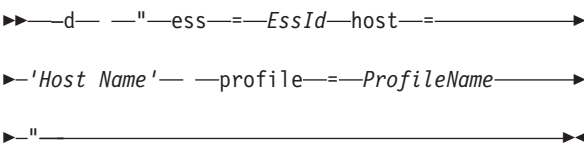
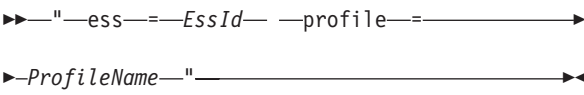
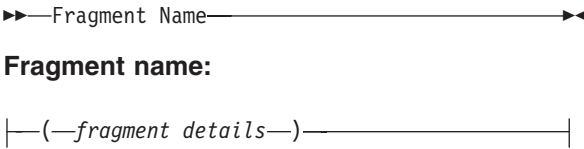
- Go to the feedback page on the website for the SAN Volume Controller Information Center at publib.boulder.ibm.com/infocenter/svc/ic/index.jsp?topic=/com.ibm.storage.svc.console.doc/feedback.htm. There you can use the feedback page to enter and submit comments or browse to the topic and use the feedback link in the running footer of that page to identify the topic for which you have a comment.
- Send your comments by email to starpubs@us.ibm.com. Include the following information for this publication or use suitable replacements for the publication title and form number for the publication on which you are commenting:
 - Publication title: *IBM System Storage SAN Volume Controller Command-Line Interface User's Guide*
 - Publication form number: GC27-2287-00
 - Page, table, or illustration numbers that you are commenting on
 - A detailed description of any information that should be changed

Syntax diagrams

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

The following table explains how to read the syntax diagrams that represent the command-line interface (CLI) commands. In doing so, it defines the symbols that represent the CLI command elements.

Element	Syntax	Description
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	▶▶—esscli————▶▶	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.

Element	Syntax	Description
Required keywords		Indicate the parameters or arguments that you must specify for the command. Required keywords appear on the main path line. Required keywords that cannot be used together are stacked vertically.
Optional keywords		Indicate the parameters or arguments that you can choose to specify for the command. Optional keywords appear below the main path line. Mutually exclusive optional keywords are stacked vertically.
Default value		Appears above the main path line.
Repeatable keyword or value		Represents a parameter or argument that you can specify more than once. A repeatable keyword or value is represented by an arrow returning to the left above the keyword or value.
Variable		Represents the value that you need to supply for a parameter or argument, such as a file name, user name, or password. Variables are in italics.
Space separator		Adds a blank space on the main path line to separate keywords, parameters, arguments, or variables from each other.
Quotation mark delimiters		Indicates the start and end of a parameter or argument that contains multiple values. Enclose one or more name-value pairs in a set of double quotation marks for a particular parameter or argument. If the value of a parameter or name-value pair contains a blank or white space, enclose the entire value in a set of single quotation marks.
Equal-sign operator		Separates a name from its value in a name-value pair.
Syntax fragment		Breaks up syntax diagrams that are too long, too complex, or repetitious. The fragment name is inserted in the main diagram, and the actual fragment is shown below the main diagram.

Terminology

These are abbreviations that are most commonly used for the command-line interface operations.

Name	Object type
Host	host
Virtual disk (volume)	vdisk
Virtual disk copy	vdiskcopy
Space-efficient (thin-provisioned) virtual disk copy	sevdiskcopy
Managed disk	mdisk
Managed disk group (storage pool)	mdiskgrp
I/O group	iogrp
Node	node
Cluster	cluster
Controller	controller
IBM FlashCopy mapping	fcmap
FlashCopy consistency group	fcconsistgrp
Metro Mirror or Global Mirror relationship	rrelationship
Metro Mirror or Global Mirror consistency group	rcconsistgrp
Unsupported/unknown object	unknown

CLI special characters

The following special characters are used in the command-line interface (CLI) command examples.

minus (-) 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. The - character cannot be used as the first character of an object name.

vertical bar (|)

A vertical bar signifies that you choose only one value. For example, [a | b] in brackets indicates that you can choose a, b, or nothing. Similarly, { a | b } in braces indicates that you must choose either a or b.

Using wildcards in the SAN Volume Controller CLI

You can use wildcards in the SAN Volume Controller Command-Line Interface.

The SAN Volume Controller supports the use of the asterisk character (*) as a wildcard within the arguments of certain parameters. There are some behavioral issues that must be considered when using wildcards in order to prevent unexpected results. These behavioral issues and the ways to avoid them are as follows:

1. Running the command while logged onto the node.

The shell will attempt to interpret any of the special characters if they are not escaped (preceded with a backslash character). Wildcards will be expanded into a list of files if any files exist that match the wildcards. If no matching files exist, the wildcard is passed to the SAN Volume Controller command untouched.

To prevent expansion, issue the following command in one of its formats:

```
svctask cleardumps -prefix '/dumps/*.txt' with single quotation marks (' '), or
```

`svctask cleardumps -prefix /dumps/*.txt` using a backslash (\), or

`svctask cleardumps -prefix "/dumps/*.txt"` with double quotation marks ("").

2. Running the command through Secure Shell (SSH), for example from a host.

This method is slightly more complicated because the host shell processes the command line before it is passed through SSH to the shell on the cluster. This means an extra layer of protection is required around the wildcard as the host shell will strip off any protecting quotes, and if the wildcard is exposed to the cluster shell, this will result in the wildcard being expanded in the cluster shell.

To prevent expansion, issue the following command in one of its formats:

`svctask cleardumps "'/dumps/*.txt'"` with single quotation marks (') inside of double quotation marks (""), or

`svctask cleardumps '/dumps/*.txt'` using a backslash (\) inside of single quotation marks ('), or

`svctask cleardumps ""/dumps/*.txt""` with double quotation marks (") inside of single quotation marks (').

Data types and value ranges

The maximum length of any single parameter that is entered into the command line is 2176 bytes.

Note: If you do not specify a name when you are creating a new object, the cluster assigns a default name. This name is made from the object type as the prefix and the object ID as the suffix. For example, a new virtual disk (VDisk) is created with ID 5. This object is given the default name of **vdisk5**. Because the system assigns these names, it does not allow you to create an object and call it **vdiskx** where *x* is the integer. This is because the cluster reserves these names (for example, `object_type_prefix integer`) for default.

Data types	Value ranges
<code>filename_arg</code>	<p>This is a (optionally fully qualified) file name. The maximum length is 169 characters. Valid characters consist of the following options:</p> <ul style="list-style-type: none">• .• /• -• _• a - z• A - Z• 0 - 9 <p>The field must not contain two consecutive '.', or start with a '.', or end with a '.'.</p>

Data types	Value ranges
directory_or_file_filter	<p>Specifies a directory and file name filter, or both, within the specified directory. Valid directory values consist of the following options:</p> <ul style="list-style-type: none"> • /dumps • /dumps/audit • /dumps/configs • /dumps/elogs • /dumps/feature • /dumps/iostats • /dumps/iotrace • /dumps/software <p>The file name filter can be any valid file name with or without the wildcard '*'. The file name filter can be appended to the end of one of the previous directory values. The maximum length is 128 characters. Valid characters consist of the following options:</p> <ul style="list-style-type: none"> • * • . • / • - • _ • a - z • A - Z • 0 - 9 <p>The field must not contain two consecutive '.', or start with a '.', or end with a '.</p>
filename_prefix	<p>This is a prefix to be used when naming a file. The maximum length is 128 characters. Valid characters consist of the following options:</p> <ul style="list-style-type: none"> • a - z • A - Z • 0 - 9 • - • _

Data types	Value ranges
name_arg	<p>Names can be specified or changed using the create and modify functions. With the view commands you can see both the name and ID of an object.</p> <p>A string of 1 - 15 characters can be used, composed of characters A - Z, a - z, 0 - 9, - (dash), and _ (underscore).</p> <p>The first character of a <i>name_arg</i> must not be numeric. The first character of an object name cannot be a - (dash) because the CLI interprets it as being the next parameter.</p> <p>When creating a name for an object, the name must not consist of the object type followed only by an integer. The exception is Metro or Global Mirror relationships, which can be named anything as long as the names are unique across the two clusters. This naming convention is used by the system to generate default names. You cannot use one of the following reserved words followed by an integer:</p> <ul style="list-style-type: none"> • cluster • controller • fccstgrp • fcmap • host • io_grp • mdisk • mdiskgrp • node • rccstgrp • rcmap <p>The cluster name is set when the cluster is created.</p>
password	<p>This is a user-defined password. A password must meet the following requirements:</p> <ul style="list-style-type: none"> • Can use a - z, A - Z, 0 - 9 in any sequence • Can use - (dash) but not as the first character • Can use _ (underscore) • Can contain a maximum of 15 characters
serial_number	<p>The format of this number conforms to IBM standard C-S 1-1121-018 1999-06 Serial Numbering for IBM products. The serial number is 7 digits, the first two of which define the manufacturing location, leaving 5 digits for the product. The standard defines a way to extend the serial number using letters in the place of numbers in the 5-digit field.</p>

Data types	Value ranges
ip_address_arg	<p>The decimal, dotted-quad-notation standard rules. The following Internet Protocol 4 (IPv4) and Internet Protocol 6 (IPv6) address formats are supported:</p> <p>IPv4 (no port set, SAN Volume Controller uses default) 1.2.3.4</p> <p>IPv4 with specific port 1.2.3.4:22</p> <p>Full IPv6, default port 1234:1234:0001:0123:1234:1234:1234:1234</p> <p>Full IPv6, default port, leading zeros suppressed 1234:1234:1:123:1234:1234:1234:1234</p> <p>Full IPv6 with port [2002:914:fc12:848:209:6bff:fe8c:4ff6]:23</p> <p>Zero-compressed IPv6, default port 2002::4ff6</p> <p>Zero-compressed IPv6 with port [2002::4ff6]:23</p>
dns_name	Dotted domain name for the subnet that the cluster is in. For example, ibm.com.
hostname	<p>The host name that is assigned to the cluster. This can be different from the cluster name and you can change the host name at any time.</p> <p>A combination of the host name and the dns_name that is used to access the cluster, for example:</p> <p>https://hostname.ibm.com/</p>
capacity_value	<p>A value with a range of 512 bytes up to 2 petabytes (PB).</p> <p>Note: The capacity can be specified as megabytes (MB), kilobytes (KB), gigabytes (GB), or PB. When MB is used, the value is specified in multiples of 512 bytes. A capacity of 0 is valid for a striped or sequential VDisk. The smallest number of supported bytes is 512.</p>
node_id	<p>A node ID differs from other IDs; a node ID is a unique ID that is assigned when a node is used to create a cluster, or when a node is added to a cluster. A <i>node_id</i> value is never reused in a cluster. Node IDs are internally represented as 64-bit numbers.</p> <p>Node IDs, like other IDs, cannot be modified by user commands.</p>

Data types	Value ranges
xxx_id	<p>All objects are referred to by unique integer IDs that are assigned by the system when the objects are created. All IDs are represented internally as 32-bit integers. Node IDs are an exception.</p> <p>IDs in the following ranges are used to identify the various types of objects:</p> <ul style="list-style-type: none"> • node_id: A positive decimal integer greater than or equal to 1 • mdisk_grp_id: 0 - 127 • io_grp_id: 0 - 3 (See Note.) • mdisk_id: 0 - 4095 • vdisk_id: 0 - 8191 • copy_id: 0 - 1 • host_id: 0 - 1023 • flash_const_grp_id: 0 - 255 • remote_const_grp_id: 0 - 255 • fcmapi_id: 0 - 4095 • rcrel_id: 0 - 8191 • controller_id: 0 - 63 <p>Note: The io_group 4 exists but is used only in certain error recovery procedures.</p> <p>These IDs, like node IDs, cannot be modified by user commands.</p> <p>Note: IDs are assigned at run time by the system and cannot be relied upon to be the same after, for example, the configuration restoration. Use object names in preference to IDs when you are working with objects.</p>
xxx_list	A colon-delimited list of values of type <i>xxx</i> .
wwpn_arg	<p>The fibre-channel worldwide port name (WWPN). This is expressed as a 64-bit hexadecimal number, for example:</p> <p>1A2B30C67AFFE47B</p> <p>These numbers must consist of the characters 0 - 9, a - f, and A - F. A command fails if you enter WWPN 0 in the command string.</p>
panel_name	A string of up to six characters that correspond to the number on the printed label below the display on the front panel of a node in the cluster.
sequence_number	A 32-bit unsigned integer, expressed in decimal format.
csi_num_arg	A 32-bit unsigned integer, expressed in decimal format.
percentage_arg	An 8-bit unsigned integer, expressed in decimal 0 - 100 format.
extent_arg	A 32-bit unsigned integer, expressed in decimal format.
num_extents_arg	A 32-bit unsigned integer, expressed in decimal format.
threads_arg	An 8-bit unsigned integer, expressed in decimal format. Valid values are 1, 2, 3, or 4.
velocity_arg	The fabric speed in gigabytes per second (GBps). Valid values are 1 or 2.
timezone_arg	The ID as detailed in the output of the svcinfo lstimezones command.
timeout_arg	The command timeout period. An integer from 0 to 600 (seconds).
stats_time_arg	The frequency at which statistics are gathered. Valid values are 1 to 60 minutes in increments of 1 minute.

Data types	Value ranges
directory_arg	<p>Specifies a directory and or file name filter within the specified directory. The following directory values are valid:</p> <ul style="list-style-type: none"> • /dumps • /dumps/audit • /dumps/cimom • /dumps/configs • /dumps/elogs • /dumps/feature • /dumps/iostats • /dumps/iotrace • /home/admin/upgrade <p>The file name filter can be any valid file name with or without the wildcard '*'. The file name filter can be appended to the end of one of the previous directory values.</p>
locale_arg	<p>The cluster locale setting. Valid values are the following: .</p> <ul style="list-style-type: none"> • 0 en_US: US English (default) • 1 zh_CN: Simplified Chinese • 2 zh_TW: Traditional Chinese • 3 ja_JP: Japanese • 4 fr_FR: French • 5 de_DE: German • 6 it_IT: Italian • 7 es_ES: Spanish
key_arg	A user-defined identifier for an SSH key. It consists of a string of up to 30 characters.
user_arg	Specifies the user: either admin or service.
copy_rate	A numeric value of 0 - 100.
copy_type	Specifies the Mirror copy type: Metro or Global.

The maximum number of values that can be entered into a colon-separated list is 128. If more than 128 items are entered into a list, an error is returned.

CLI commands and parameters

CLI commands and parameters are represented in the syntax diagram.

The SAN Volume Controller command-line interface offers command line completion for command entry. Command line completion allows you to type in the first few characters of a command and press the Tab key to fill in the rest of the command name. If there are multiple commands that start with the same characters, then a list of possible commands is returned. You can type in more characters until the command name is unambiguous.

CLI parameters can be entered in any order except in the following situations:

- The first argument following the command name must be the action that is to be performed.
- Where you are performing an action against a specific object, the object ID or name must be the last argument in the line.

A valid parameter meets the following requirements:

- Parameters can be entered in any order.

- If a parameter has an associated argument, the argument must *always* follow the parameter.
- A parameter *must* start with a '-'; otherwise, it is assumed to be an argument.
- The maximum length of any single parameter that can be entered into the CLI is 128 bytes.
- An argument can contain multiple data items. The maximum number of data items that you can enter into such a list is 128. For a component list, separate the individual items by a colon.
- Any parameter with an argument can be entered as -parameter=argument.
- Entering -param= means the argument is an empty string, equivalent to -param.
- The symbol '--' is valid as the next to last entry on the command line. It specifies that the next entry is the target object name or ID, even if it begins with a hyphen.
svctask chuser -usergrp=-usergrp -- -password
- The symbol '-' is valid as the final word on the command line.

Examples that are valid:

```
svctask mkuser -name fred -usergrp 0 -password buckets
svctask mkuser -name fred -usergrp 0 -password=buckets
svctask mkuser -name fred -usergrp 0 -password=buckets --
svctask mkuser -name=-barney -usergrp=0 -password=buckets
```

```
svctask chuser -usergrp 1 fred
svctask chuser -usergrp 1 -- fred
svctask chuser -usergrp 1 -- -barney
```

Examples that are invalid:

```
svctask chuser -usergrp 1 fred --
svctask chuser -usergrp 1 -- fred --
svctask chuser -- -usergrp 1 fred
svctask chuser -usergrp 1 -barney
```

CLI flags

The following flags are common to all command-line interface (CLI) commands.

-? or -h

Print help text. For example, issuing **svcinfolcluster -h** provides a list of the actions available with the **svcinfolcluster** command.

-nomsg

When used, this flag prevents the display of the successfully created output. For example, if you issue the following command:

```
svctask mkmdiskgrp -ext 16
```

it displays:

```
MDisk Group, id [6], successfully created
```

However, if the -nomsg parameter is added, for example:

```
svctask mkmdiskgrp -ext 16 -nomsg
```

the following information is displayed:

```
6
```

This parameter can be entered for any command, but is only acted upon by those commands that generate the successfully created outputs. All other commands ignore this parameter.

CLI messages

Ensure that you are familiar with the command-line interface (CLI) messages.

When some commands complete successfully, textual output is normally provided. However, some commands do not provide any output. The phrase `No feedback` is used to indicate that no output is provided. If the command does not complete successfully, an error is generated. For example, if the command has failed as a result of the cluster being unstable, the following output is provided:

- `CMMVC5786E` The action failed because the cluster is not in a stable state.

Chapter 1. Preparing the SSH client

Secure Shell (SSH) is a client-server network application. It is a communication vehicle between the host system and the SAN Volume Controller command-line interface (CLI).

Overview

The SAN Volume Controller cluster acts as the SSH server in this relationship. The SSH client provides a secure environment in which to connect to a remote machine. It uses the principles of public and private keys for authentication.

SSH keys are generated by the SSH software. This includes a public key, which is uploaded and maintained by the cluster, and a private key that is kept private to the host that is running the SSH client. These keys authorize specific users to access the administration and service functions on the cluster. Each key is associated with a user on the cluster. Up to 400 users can be defined on the cluster. You can also create new users and assign keys to them.

Authenticating SSH logins

When you are using AIX® hosts, SSH logins are authenticated on the cluster using the RSA-based authentication that is supported in the OpenSSH client that is available for AIX. This scheme is based on public-key cryptography, using an algorithm known commonly as RSA.

Note: The authentication process for non-AIX hosts systems is similar.

With this scheme (as in similar OpenSSH systems on other host types), the encryption and decryption is done using separate keys. This means that it is not possible to derive the decryption key from the encryption key.

Because physical possession of the private key allows access to the cluster, the private key must be kept in a protected place, such as the .ssh directory on the AIX host, with restricted access permissions.

When SSH client (A) attempts to connect to SSH server (B), the key pair authenticates the connection. The key consists of two halves: the public keys and private keys. The SSH client public key is put onto SSH Server (B) using some means outside of the SSH session. When SSH client (A) tries to connect, the private key on SSH client (A) is able to authenticate with its public half on SSH server (B).

To connect to the cluster, the SSH client requires a user login name and a key pair. When using an SSH client to access a SAN Volume Controller cluster, you must always specify a user login name of admin. The SAN Volume Controller cluster uses the key pair to identify the user accessing the cluster.

Preparing the SSH client on a Windows host

- | If you use the Secure Shell (SSH) to log in to the SAN Volume Controller or Storwize® V7000, you must use the RSA-based private key authentication, instead of a typical user name and password. This topic describes how to prepare the SSH client on a Windows host.
- | The IBM System Storage Productivity Center (SSPC) and the workstation for the SAN Volume Controller include the PuTTY client program, which is a Microsoft Windows SSH client program. The PuTTY client program can be installed on your SSPC or workstation server in one of the following ways:
 - | • If you purchased the SSPC or the workstation hardware option from IBM, the PuTTY client program has been preinstalled on the hardware.

- You can use the workstation software installation CD to install the PuTTY client program. The SSPC, workstation hardware option, and the software-only workstation each provide this CD.
- You can use the separate PuTTY client program-installation wizard, **putty-version-installer.exe**. You can download the PuTTY client program from the following Web site:
www.chiark.greenend.org.uk/~sgtatham/putty/

Note: Before you install the PuTTY client program, ensure that your Windows system meets the system requirements. See the *IBM System Storage Productivity Center Introduction and Planning Guide* for system requirements.

If you want to use an SSH client other than the PuTTY client, the following Web site offers SSH client alternatives for Windows:

www.openssh.org/windows.html

Generating an SSH key pair using PuTTY

You must generate a Secure Shell (SSH) key pair to use the SAN Volume Controller command-line interface (CLI).

Perform the following steps to generate SSH keys using the PuTTY key generator (PuTTYgen):

1. Start PuTTYgen by clicking **Start** → **Programs** → **PuTTY** → **PuTTYgen**. The PuTTY Key Generator panel is displayed.
2. Click **SSH-2 RSA** as the type of key to generate.

Note: Leave the number of bits in a generated key value at 1024.

3. Click **Generate** and then move the cursor around the blank area of the Key section to generate the random characters that create a unique key. When the key has been completely generated, the information about the new key is displayed in the Key section.

Attention: Do not modify the Key fingerprint or the Key comment fields; this can cause your key to no longer be valid.

4. (Optional) Enter a passphrase in the **Key passphrase** and **Confirm passphrase** fields. The passphrase encrypts the key on the disk; therefore, it is not possible to use the key without first entering the passphrase.
5. Save the public key by performing the following steps:
 - a. Click **Save public key**. You are prompted for the name and location of the public key.
 - b. Type `icat.pub` as the name of the public key and specify the location where you want to save the public key. For example, you can create a directory on your computer called *keys* to store both the public and private keys.
 - c. Click **Save**.
6. Save the private key by performing the following steps:
 - a. Click **Save private key**. The PuTTYgen Warning panel is displayed.
 - b. Click **Yes** to save the private key without a passphrase.
 - c. Type `icat` as the name of the private key, and specify the location where you want to save the private key. For example, you can create a directory on your computer called *keys* to store both the public and private keys. It is recommended that you save your public and private keys in the same location.
 - d. Click **Save**.
7. Close the PuTTY Key Generator window.

Configuring a PuTTY session for the CLI

You must configure a PuTTY session using the Secure Shell (SSH) key pair that you have generated before you can use the command-line interface (CLI).

Attention: Do not run scripts that create child processes that run in the background and call SAN Volume Controller commands. This can cause the system to lose access to data and cause data to be lost.

Perform the following steps to configure a PuTTY session for the CLI:

1. Select **Start** → **Programs** → **PuTTY** → **PuTTY**. The PuTTY Configuration window opens.
2. Click **Session** in the Category navigation tree. The Basic options for your PuTTY session are displayed.
3. Click **SSH** as the Protocol option.
4. Click **Only on clean exit** as the Close window on exit option. This ensures that connection errors are displayed.
5. Click **Connection** → **SSH** in the Category navigation tree. The options controlling SSH connections are displayed.
6. Click **2** as the Preferred SSH protocol version.
7. Click **Connection** → **SSH** → **Auth** in the Category navigation tree. The Options controller SSH authentication are displayed.
8. Click **Browse** or type the fully qualified file name and location of the SSH client and private key in the **Private key file for authentication** field.
9. Click **Connection** → **Data** in the Category navigation tree.
10. Type **admin** in the **Auto-login username** field.
11. Click **Session** in the Category navigation tree. The Basic options for your PuTTY session are displayed.
12. In the **Host Name (or IP Address)** field, type the name or IP address of one of the SAN Volume Controller cluster IP addresses or host names.
13. Type **22** in the **Port** field. The SAN Volume Controller cluster uses the standard SSH port.
14. Type the name that you want to use to associate with this session in the **Saved Sessions** field. For example, you can name the session **SAN Volume Controller Cluster 1**.
15. Click **Save**.

You have now configured a PuTTY session for the CLI.

Note: If you configured more than one IP address for the SAN Volume Controller cluster, repeat the previous steps to create another saved session for the second IP address. This can then be used if the first IP address is unavailable.

Connecting to the CLI using PuTTY

Ensure that you are familiar with how to run the PuTTY and plink utilities.

Note: Windows users can download PuTTY from the following Web site: <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>(<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>).

The Secure Shell (SSH) protocol specifies that the first access to a new host server sends a challenge to the SSH user to accept the SSH server public key. Because this is the first time that you connect to an SSH server, the server is not included in the SSH client list of known hosts. Therefore, there is a fingerprint challenge, which asks if you accept the responsibility of connecting with this host. If you type **y**, the host fingerprint and IP address are saved by the SSH client.

| When you use PuTTY, you must also type *y* to accept this host fingerprint. However, the host fingerprint and IP address are stored in the registry for the user name that is logged onto Windows.

| The SSH protocol also specifies that once the SSH server public key is accepted, another challenge is presented if the fingerprint of an SSH server changes from the one previously accepted. In this case, you must decide if you want to accept this changed host fingerprint.

| **Note:** The SSH server keys on the SAN Volume Controller are regenerated when a microcode load is performed on the cluster. As a result, a challenge is sent because the fingerprint of the SSH server has changed.

| All command-line interface (CLI) commands are run in an SSH session. You can run the commands in one of the following modes:

- | • An interactive prompt mode
- | • A single line command mode, which is entered one time to include all parameters

| **Interactive mode**

| For interactive mode, you can use the PuTTY executable to open the SSH restricted shell.

| The following is an example of the command that you can issue to start interactive mode:

| `C:\support_utils\putty admin@svconsoleip`

| where `support_utils\putty` is the location of your `putty.exe` file and `svconsoleip` is the IP address of your management GUI.

| If you were to issue the `svcinfolssshkeys` command, which lists the SSH client public keys that are stored on the SAN Volume Controller cluster, the following output is displayed:

```
IBM_2145:your_cluster_name:admin>svcinfolssshkeys -user all -delim :
id:userid:key identifier
1:admin:smith
2:admin:jones
```

| You can type `exit` and press **Enter** to escape the interactive mode command.

| The following is an example of the host fingerprint challenge when using `plink` in interactive mode:

```
C:\Program Files\IBM\svconsole\cimom>plink admin@9.43.225.208
The server's host key is not cached in the registry. You
have no guarantee that the server is the computer you
think it is.
The server's key fingerprint is:
ssh-rsa 1024 e4:c9:51:50:61:63:e9:cd:73:2a:60:6b:f0:be:25:bf
If you trust this host, enter "y" to add the key to
PuTTY's cache and carry on connecting.
If you want to carry on connecting just once, without
adding the key to the cache, enter "n".
If you do not trust this host, press Return to abandon the
connection.
Store key in cache? (y/n) y
Using username "admin".
Authenticating with public key "imported-openssh-key"
IBM_2145:your_cluster_name:admin>
```

Single line command

For single line command mode, you can type the following all on one command line:

```
C:\Program Files\IBM\svccconsole\cimom>
plink admin@9.43.225.208 svcinfo lssshkeys
-user all -delim :
Authenticating with public key "imported-openssh-key"
id:userid:key identifier
1:admin:smith
2:admin:jones
```

Note: If you are submitting a CLI command with all parameters in single line command mode, you are challenged upon first appearance of the SSH server host fingerprint. Ensure that the SSH server host fingerprint is accepted before you submit a batch script file.

The following is an example of the host fingerprint challenge when using plink in single line command mode:

```
C:\Program Files\IBM\svccconsole\cimom>
plink admin@9.43.225.208 svcinfo lssshkeys
-user all -delim :
The server's host key is not cached in the registry. You
have no guarantee that the server is the computer you
think it is.
The server's key fingerprint is:
ssh-rsa 1024 e4:c9:51:50:61:63:e9:cd:73:2a:60:6b:f0:be:25:bf
If you trust this host, enter "y" to add the key to
PuTTY's cache and carry on connecting.
If you want to carry on connecting just once, without
adding the key to the cache, enter "n".
If you do not trust this host, press Return to abandon the
connection.
Store key in cache? (y/n) y
Authenticating with public key "imported-openssh-key"
/bin/ls: /proc/20282/exe: Permission denied
dircolors: ~/etc/DIR_COLORS': Permission denied
id:userid:key identifier
1:admin:smith
2:admin:jones
```

Starting a PuTTY session for the CLI

You must start a PuTTY session to connect to the command-line interface (CLI).

This task assumes that you have already configured and saved a PuTTY session using the Secure Shell (SSH) key pair that you created for the CLI.

Perform the following steps to start a PuTTY session:

1. Select **Start** → **Programs** → **PuTTY** → **PuTTY**. The PuTTY Configuration window opens.
2. Select the name of your saved PuTTY session and click **Load**.
3. Click **Open**.

Note: If this is the first time that the PuTTY application is being used since you generated and uploaded the SSH key pair, a PuTTY Security Alert window is displayed. Click **Yes** to accept the change and trust the new key.

4. Type **admin** in the **login as:** field and press Enter.

Preparing the SSH client on an AIX or Linux host

If you use the Secure Shell (SSH) to log in to the SAN Volume Controller or Storwize® V7000, you must use the RSA-based private key authentication, instead of a typical user name and password. This topic describes how to prepare the SSH client on an AIX or Linux host.

Ensure that you have an SSH client installed on your system:

IBM AIX operating systems

For IBM AIX 5L™ for POWER®, versions 5.1, 5.2, 5.3, and AIX version 6.1 for IBM POWER6® architecture, you can obtain the OpenSSH client from the bonus packs, but you also must obtain its prerequisite, OpenSSL, from the IBM AIX toolbox for Linux applications for IBM Power Systems™. For AIX 4.3.3, you can obtain the software from the AIX toolbox for Linux applications. You can also obtain the AIX installation images from IBM developerWorks® at the following Web site:

oss.software.ibm.com/developerworks/projects/openssh

Linux operating systems

The OpenSSH client is installed by default on most Linux distributions. If it is not installed on your system, consult your Linux installation documentation or visit the following Web site:

www.openssh.org/portable.html

The OpenSSH client can run on a variety of additional operating systems. For more information about the openSSH client, visit the following Web site:

www.openssh.org/portable.html

Perform the following steps to set up an RSA key pair on the AIX or Linux host and the SAN Volume Controller or Storwize® V7000 cluster:

1. Create an RSA key pair by issuing a command on the host that is similar to the following command:

```
ssh-keygen -t rsa
```

Tip: Issue the command from the \$HOME/.ssh directory.

This process generates two user named files. If you select the name *key*, the files are named *key* and *key.pub*. Where *key* is the name of the private key and *key.pub* is the name of the public key.

2. Associate the public key with a user on the SAN Volume Controller or Storwize® V7000 system. using the management GUI.

You can now access the cluster with an SSH command similar to the following:

```
ssh -i full_path_to_key admin@my_cluster
```

Where *my_cluster* is the name of the cluster IP, and *full_path_to_key* is the full path to the *key* file that was generated in the previous step. Always use *admin* as the SSH user name. The SAN Volume Controller software determines which user is logging in from the key they are using.

Note: You can omit `-i full_path_to_key` if you configure the SSH client to use the key file automatically. For more information, refer to the OpenSSH documentation.

Generating an SSH key pair using OpenSSH

This topic describes how to generate an SSH key pair using OpenSSH.

Perform the following steps to set up an RSA key pair on the AIX or Linux host and the SAN Volume Controller or Storwize® V7000 cluster:

1. Create an RSA key pair by issuing a command on the host that is similar to the following command:

```
ssh-keygen -t rsa
```

Tip: Issue the command from the `$HOME/.ssh` directory. This process generates two user named files. If you select the name *key*, the files are named *key* and *key.pub*. Where *key* is the name of the private key and *key.pub* is the name of the public key.
2. Associate the public key with a user on the SAN Volume Controller or Storwize® V7000 system. using the management GUI.

Connecting to the CLI using OpenSSH

This topic describes how to connect to the CLI using OpenSSH.

You can access the cluster with an SSH command similar to the following:

```
ssh -i full_path_to_key admin@my_cluster
```

Where *my_cluster* is the name of the cluster IP, and *full_path_to_key* is the full path to the *key* file that was generated in the previous step. Always use *admin* as the SSH user name. The SAN Volume Controller software determines which user is logging in from the key they are using.

Note: You can omit `-i full_path_to_key` if you configure the SSH client to use the key file automatically. For more information, refer to the OpenSSH documentation.

Creating users

You can create either a local or a remote user to access a SAN Volume Controller cluster.

You can create two categories of users that access the cluster. These types are based on how the users are authenticated to the cluster. Local users must provide either a password, a Secure Shell (SSH) key, or both. Local users are authenticated through the authentication methods that are located on the SAN Volume Controller cluster. If the local user needs access to management GUI, a password is needed for the user. If the user requires access to the command-line interface (CLI) then a valid SSH key file is necessary. If a user is working with both interfaces, then both a password and SSH key are required. Local users must be part of a user group that is defined on the cluster. User groups define roles that authorize the users within that group to a specific set of operations on the cluster.

A remote user is authenticated on a remote service usually provided by a SAN management application, such as IBM Tivoli® Storage Productivity Center, and does not need local authentication methods. For a remote user, both a password and SSH key are required to use the command-line interface. Remote users only need local credentials to access to the management GUI if the remote service is down. Remote users have their groups defined by the remote authentication service.

Complete the following steps to create either a local or remote user:

1. Select **User Management** → **Users** .
2. Click **New User** .
3. Enter the information on the new user and click **OK**.

Chapter 2. Copying the SAN Volume Controller software upgrade files using PuTTY scp

PuTTY scp (pscp) provides a file transfer application for secure shell (SSH) to copy files either between two directories on the configuration node or between the configuration node and another host.

To use the pscp application, you must have the appropriate permissions on the source and destination directories on your respective hosts.

The pscp application is available when you install an SSH client on your host system. You can access the pscp application through a Microsoft Windows command prompt.

Perform the following steps to use the pscp application:

1. Start a PuTTY session.
2. Configure your PuTTY session to access your SAN Volume Controller cluster.
3. Save your PuTTY configuration session. For example, you can name your saved session SVCPUTTY.
4. Open a command prompt.
5. Issue the following command to set the path environment variable to include the PuTTY directory:

```
set path=C:\Program Files\putty;%path%
```

where *Program Files* is the directory where PuTTY is installed.

6. Issue the following command to copy the package onto the node where the CLI runs:

```
pscp -load saved_putty_configuration  
directory_software_upgrade_files/software_upgrade_file_name  
admin@cluster_ip_address:/home/admin/upgrade
```

where *saved_putty_configuration* is the name of the PuTTY configuration session,

directory_software_upgrade_files is the location of the software upgrade files, *software_upgrade_file_name* is the name of the software upgrade file, and *cluster_ip_address* is an IP address of your cluster.

If there is insufficient space to store the software upgrade file on the cluster, the copy process fails.

Perform the following steps:

- a. Use pscp to copy data that you want to preserve from the /dumps directory. For example, issue the following command to copy all event logs from the cluster to the IBM System Storage Productivity Center or the master console:

```
pscp -unsafe -load saved_putty_configuration  
admin@cluster_ip_address:/dumps/e/logs/*  
your_preferred_directory
```

where *saved_putty_configuration* is the name of the PuTTY configuration session, *cluster_ip_address* is the IP address of your cluster, and *your_preferred_directory* is the directory where you want to transfer the event logs.

- b. Issue the **svctask cleardumps** command to free space on the cluster:

```
svctask cleardumps -prefix /dumps
```

- c. Then repeat step 6.

Chapter 3. Using the CLI

The SAN Volume Controller cluster command-line interface (CLI) is a collection of commands that you can use to manage the SAN Volume Controller.

Overview

The CLI commands use the Secure Shell (SSH) connection between the SSH client software on the host system and the SSH server on the SAN Volume Controller cluster.

Before you can use the CLI, you must have already created a cluster.

You must perform the following actions to use the CLI from a client system:

- Install and set up SSH client software on each system that you plan to use to access the CLI.
- Generate an SSH key pair on each SSH client.
- Store the SSH public key for each SSH client on the SAN Volume Controller.

Note: After the first SSH public key is stored, you can add additional SSH public keys using either the management GUI or the CLI.

You can use the CLI to perform the following functions:

- Set up of the cluster, its nodes, and the I/O groups
- | • Analyze error logs (event logs)
- | • Set up and maintenance of managed disks (MDisk) and MDisk groups (storage pools)
- Set up and maintenance of client public SSH keys on the cluster
- Set up and maintenance of virtual disks (VDisks)
- Set up of logical host objects
- | • Map VDisks (volumes) to hosts
- Navigate from managed hosts to VDisks and to MDisks, and the reverse direction up the chain
- Set up and start Copy Services:
 - FlashCopy and FlashCopy consistency groups
 - Synchronous Metro Mirror and Metro Mirror consistency groups
 - Asynchronous Global Mirror and Global Mirror consistency groups

Setting the cluster time using the CLI

You can use the command-line interface (CLI) to set the cluster time.

Perform the following steps to set the cluster time:

1. Issue the **svcinfolistimezone** CLI command to display the current time-zone settings for the cluster. The time zone and the associated time-zone ID are displayed.
2. Issue the **svcinfolistimezones** CLI command to list the time zones that are available on the cluster. A list of valid time-zone settings are displayed. Each time zone is assigned an ID. The time zone and the associated ID are indicated in the list.
3. Issue the following CLI command to set the time zone for the cluster.

```
svctask settimezone -timezone time_zone_setting
```

where *time_zone_setting* is the new time zone ID that you have chosen from the list of time zones that are available on the cluster.

4. Issue the following CLI command to set the time for the cluster:

```
svctask setclustertime -time 031809142005
```

where *031809142005* is the new time that you want to set for the cluster. You must use the *MMDDHHmmYYYY* format to set the time for the cluster.

Viewing and updating license settings using the CLI

You can use the command-line interface (CLI) to view and update your license settings.

SAN Volume Controller provides two license options: Physical Disk Licensing and Capacity Licensing. Perform the following steps to view and update your SAN Volume Controller license settings:

1. Issue the **svcinfolicense** CLI command to view the current license settings for the cluster.
2. Issue the **svctask chlicense** CLI command to change the licensed settings of the cluster.

Attention:

- License settings are entered when the cluster is first created; do not update the settings unless you have changed your license.
- To select Physical Disk Licensing, run the **svctask chlicense** command with one or more of the **physical_disks**, **physical_flash**, and **physical_remote** parameters.
- To select Capacity Licensing, run the **svctask chlicense** command with one or more of the **-flash**, **-remote**, and **-virtualization** parameters. If the physical disks value is nonzero, these parameters cannot be set.

Displaying cluster properties using the CLI

You can use the command-line interface (CLI) to display the properties for a cluster.

Perform the following step to display cluster properties:

Issue the **svcinfolcluster** command to display the properties for a cluster.

The following is an example of the command you can issue:

```
| svcinfolcluster -delim : build1
```

where *build1* is the name of the cluster.

```

id:000002007A00A0FE
name:build1
location:local
partnership:
bandwidth:
total_mdisk_capacity:90.7GB
space_in_mdisk_grps:90.7GB
space_allocated_to_vdisks:14.99GB
total_free_space:75.7GB
statistics_status:on
statistics_frequency:15
required_memory:0
cluster_locale:en_US
time_zone:522 UTC
code_level:6.1.0.0 (build 47.3.1009031000)
FC_port_speed:2Gb
console_IP:9.71.46.186:443
id_alias:000002007A00A0FE
gm_link_tolerance:300
gm_inter_cluster_delay_simulation:0
gm_intra_cluster_delay_simulation:0
email_reply:
email_contact:
email_contact_primary:
email_contact_alternate:
email_contact_location:
email_state:stopped
inventory_mail_interval:0
total_vdiskcopy_capacity:15.71GB
total_used_capacity:13.78GB
total_overallocation:17
total_vdisk_capacity:11.72GB
cluster_ntp_IP_address:
cluster_isns_IP_address:
iscsi_auth_method:none
iscsi_chap_secret:
auth_service_configured:no
auth_service_enabled:no
auth_service_url:
auth_service_user_name:
auth_service_pwd_set:no
auth_service_cert_set:no
relationship_bandwidth_limit:25
gm_max_host_delay:5
tier:generic_ssd
tier_capacity:0.00MB
tier_free_capacity:0.00MB
tier:generic_hdd
tier_capacity:90.67GB
tier_free_capacity:75.34GB
email_contact2:
email_contact2_primary:
email_contact2_alternate:
total_allocated_extent_capacity:16.12GB

```

Maintaining passwords for the front panel using the CLI

You can use the command-line interface (CLI) to view and change the status of the password reset feature for the SAN Volume Controller front panel.

The cluster superuser password can be reset using the front panel of the configuration node. To meet varying security requirements, this functionality can be enabled or disabled using the CLI.

Complete the following steps to view and change the status of the password reset feature:

1. Issue the **svctask setpwdreset** CLI command to view and change the status of the password reset feature for the SAN Volume Controller front panel.
2. Record the cluster superuser password because you cannot access the cluster without it.

Storwize® V7000: The cluster superuser password can be reset using a USB key. To meet varying security requirements, this functionality can be enabled or disabled using the CLI. Complete the following steps to view and change the status of the password reset feature:

1. Issue the **svctask setpwdreset** CLI command to view and change the status of the password reset feature for the Storwize® V7000.
2. Record the cluster superuser password because you cannot access the cluster without it.

Re-adding a repaired node to a cluster using the CLI

You can use the command-line interface (CLI) to re-add a failed node back into a cluster after it was repaired.

Before you add a node to a cluster, you must make sure that the switchd\ zoning is configured such that the node being added is in the same zone as all other nodes in the cluster. If you are replacing a node and the switch is zoned by worldwide port name (WWPN) rather than by switch port, make sure that the switch is configured such that the node being added is in the same VSAN/zone.

Attention:

1. If you are re-adding a node to the SAN, ensure that you are adding the node to the same I/O group from which it was removed. Failure to do this can result in data corruption. You must use the information that was recorded when the node was originally added to the cluster. If you do not have access to this information, call the IBM Support Center to add the node back into the cluster without corrupting the data.
2. The LUNs that are presented to the ports on the new node must be the same as the LUNs that are presented to the nodes that currently exist in the cluster. You must ensure that the LUNs are the same before you add the new node to the cluster.
3. LUN masking for each LUN must be identical on all nodes in a cluster. You must ensure that the LUN masking for each LUN is identical before you add the new node to the cluster.
4. You must ensure that the model type of the new node is supported by the SAN Volume Controller software level that is currently installed on the cluster. If the model type is not supported by the SAN Volume Controller software level, upgrade the cluster to a software level that supports the model type of the new node. See the following Web site for the latest supported software levels:
Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

Special procedures when adding a node to a cluster

Applications on the host systems direct I/O operations to file systems or logical volumes that are mapped by the operating system to virtual paths (vpaths), which are pseudo disk objects supported by the Subsystem Device Driver (SDD). SDD maintains an association between a VPath and a SAN Volume Controller virtual disk (VDisk). This association uses an identifier (UID) which is unique to the VDisk and is never reused. The UID allows SDD to directly associate vpaths with VDIsks.

SDD operates within a protocol stack that contains disk and fibre channel device drivers that allow it to communicate with the SAN Volume Controller using the SCSI protocol over fibre channel as defined by the ANSI FCS standard. The addressing scheme provided by these SCSI and fibre-channel device drivers uses a combination of a SCSI logical unit number (LUN) and the worldwide node name (WWNN) for the fibre channel node and ports.

If an error occurs, the error recovery procedures (ERPs) operate at various tiers in the protocol stack. Some of these ERPs cause I/O to be redriven using the same WWNN and LUN numbers that were previously used.

SDD does not check the association of the VDisk with the VPath on every I/O operation that it performs.

Before you add a node to the cluster, you must check to see if any of the following conditions are true:

- The cluster has more than one I/O group.
- The node being added to the cluster uses physical node hardware or a slot which has previously been used for a node in the cluster.
- The node being added to the cluster uses physical node hardware or a slot which has previously been used for a node in another cluster and both clusters have visibility to the same hosts and back-end storage.

If any of the previous conditions are true, the following special procedures apply:

- The node must be added to the same I/O group that it was previously in. You can use the command-line interface (CLI) command **svcinfolnode** or the management GUI to determine the WWN of the cluster nodes.
- Before you add the node back into the cluster, you must shut down all of the hosts using the cluster. The node must then be added before the hosts are rebooted. If the I/O group information is unavailable or it is inconvenient to shut down and reboot all of the hosts using the cluster, then do the following:
 - On all of the hosts connected to the cluster, unconfigure the fibre-channel adapter device driver, the disk device driver and multipathing driver before you add the node to the cluster.
 - Add the node to the cluster and then reconfigure the fibre-channel adapter device driver, the disk device driver, and multipathing driver.

Scenarios where the special procedures can apply

The following two scenarios describe situations where the special procedures can apply:

- Four nodes of an eight-node cluster have been lost because of the failure of a pair of 2145 UPS or four 2145 UPS-1U. In this case, the four nodes must be added back into the cluster using the CLI command **svctask addnode** or the management GUI.
- A user decides to delete four nodes from the cluster and add them back into the cluster using the CLI command **svctask addnode** or the management GUI.

For 5.1.0 nodes, the SAN Volume Controller automatically re-adds nodes that have failed back to the cluster. If the cluster reports an error for a node missing (error code 1195) and that node has been repaired and restarted, the cluster automatically re-adds the node back into the cluster. This process can take up to 20 minutes, so you can manually re-add the node by completing the following steps:

1. Issue the **svcinfolnode** CLI command to list the nodes that are currently part of the cluster and determine the I/O group for which to add the node.

The following is an example of the output that is displayed:

```
svcinfolnode -delim :  
  
id:name:UPS_serial_number:WWNN:status:I0_group_id:I0_group_name  
:config_node:UPS_unique_id:hardware:iscsi_name:iscsi_alias  
:panel_name:enclosure_id:canister_id:enclosure_serial_number  
1:node1::50050868010050B2:online:0:io_grp0:yes::100:iqn.1986-03.com.ibm  
:2145.cluster0.node1::02-1:2:1:123ABCG  
2:node2::50050869010050B2:online:0:io_grp0:no::100:iqn.1986-03.com.ibm  
:2145.cluster0.node2::02-2:2:2:123ABDG
```

Storwize® V7000

```
svcinfolnode -delim :
```

```
id:name:UPS_serial_number:WWNN:status:IO_group_id:IO_group_name
:config_node:UPS_unique_id:hardware:iscsi_name:iscsi_alias
:panel_name:enclosure_id:canister_id:enclosure_serial_number
1:node1::50050868010050B2:online:0:io_grp0:yes::100:iqn.1986-03.com.ibm
:2145.cluster0.node1::02-1:2:1:123ABCG
2:node2::50050869010050B2:online:0:io_grp0:no::100:iqn.1986-03.com.ibm
:2145.cluster0.node2::02-2:2:2:123ABDG
```

2. Issue the **svcinfolnodecandidate** CLI command to list nodes that are not assigned to a cluster and to verify that a second node is added to an I/O group.

The following is an example of the output that is displayed:

```
svcinfolnodecandidate -delim :
```

```
id:panel_name:UPS_serial_number:UPS_unique_id:hardware
5005076801000001:000341:10L3ASH:202378101C0D18D8:8A4
5005076801000009:000237:10L3ANF:202378101C0D1796:8A4
50050768010000F4:001245:10L3ANF:202378101C0D1796:8A4
....
```

3. Issue the **svctask addnode** CLI command to add a node to the cluster.

Important: Each node in an I/O group must be attached to a different uninterruptible power supply.

The following is an example of the CLI command you can issue to add a node to the cluster using the panel name parameter:

```
svctask addnode -panelname 000237
-iogrp io_grp0
```

Where *000237* is the panel name of the node, *io_grp0* is the name of the I/O group that you are adding the node to.

The following is an example of the CLI command you can issue to add a node to the cluster using the WWNN parameter:

```
svctask addnode -wwnodename 5005076801000001
-iogrp io_grp1
```

Where *5005076801000001* is the WWNN of the node, *io_grp1* is the name of the I/O group that you are adding the node to.

4. Issue the **svcinfolnode** CLI command to verify the final configuration.

The following example shows output that is displayed:

```
svcinfolnode -delim :
```

```
id:name:UPS_serial_number:WWNN:status:IO_group_id:IO_group_name:config_node:UPS_unique_id:
hardware:iscsi_name:iscsi_alias
1:node1:10L3ASH:0000000000000000:offline:0:io_grp0:no:1000000000003206:
8A4:iqn.1986-03.com.ibm:2145.ndihill.node1:
```

Record the following information for the new node:

- Node name
- Node serial number
- WWNN
- IQNs (if using hosts attached using iSCSI connections)
- All WWPNS
- I/O group that contains the node

Note: If this command is issued quickly after you have added nodes to the cluster, the status of the nodes might be adding. The status is shown as adding if the process of adding the nodes to the cluster is still in progress. You do not have to wait for the status of all the nodes to be online before you continue with the configuration process.

The nodes have been added to the cluster.

Displaying node properties using the CLI

You can use the command-line interface (CLI) to display node properties.

Perform the following steps to display the node properties:

1. Issue the **svcinfo lsnode** CLI command to display a concise list of nodes in the cluster.

The following is an example of the CLI command you can issue to list the nodes in the cluster:

```
svcinfo lsnode -delim :
```

The following is an example of the output that is displayed:

```
id:name:UPS_serial_number:WWNN:status:IO_group_id:IO_group_name:config_node:UPS_unique_id:hardware:iscsi_name:iscsi_alias:
panel_name:enclosure_id:canister_id:enclosure_serial_number
1:node1:UPS_Fake_SN:50050768010050B1:online:0:io_grp0:yes:10000000000050B1:8G4:iqn.1986-03.com.ibm:2145.cluster0.node1:000368:::
```

2. Issue the **svcinfo lsnode** CLI command and specify the node ID or name of the node that you want to receive detailed output.

The following is an example of the CLI command you can issue to list detailed output for a node in the cluster:

```
svcinfo lsnode -delim : group1node1
```

Where *group1node1* is the name of the node for which you want to view detailed output.

The following is an example of the output that is displayed:

```
id:1
name:group1node1
UPS_serial_number:10L3ASH
WWNN:500507680100002C
status:online
IO_group_id:0
IO_group_name:io_grp0
partner_node_id:2
partner_node_name:group1node2
config_node:yes
UPS_unique_id:202378101C0D18D8
port_id:500507680110002C
port_status:active
port_speed:2GB
port_id:500507680120002C
port_status:active
port_speed:2GB
port_id:500507680130002C
port_status:active
port_speed:2GB
port_id:500507680140003C
port_status:active
port_speed:2GB
hardware:8A4
iscsi_name:iqn.1986-03.com.ibm:2145.ndihill.node2
iscsi_alias
failover_active:no
failover_name:node1
failover_iscsi_name:iqn.1986-03.com.ibm:2145.ndihill.node1
failover_iscsi_alias
```

Discovering MDisks using the CLI

You can use the command-line interface (CLI) to discover managed disks (MDisks).

When back-end controllers are added to the fibre-channel SAN and are included in the same switch zone as a SAN Volume Controller cluster, the cluster automatically discovers the back-end controller and integrates the controller to determine the storage that is presented to the SAN Volume Controller nodes. The SCSI logical units (LUs) that are presented by the back-end controller are displayed as unmanaged MDisks. However, if the configuration of the back-end controller is modified after this has occurred, the SAN Volume Controller cluster might be unaware of these configuration changes. You can request that the SAN Volume Controller cluster rescans the fibre-channel SAN to update the list of unmanaged MDisks.

Note: The automatic discovery that is performed by SAN Volume Controller cluster does not write anything to an unmanaged MDisk. You must instruct the SAN Volume Controller cluster to add an MDisk to an MDisk group (storage pool) or use an MDisk to create an image mode virtual disk (VDisk).

Perform the following steps to discover and then view a list of MDisks:

1. Issue the **svctask detectmdisk** CLI command to manually scan the fibre-channel network. The scan discovers any new MDisks that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

Notes:

- a. Only issue the **svctask detectmdisk** command when you are sure that all of the disk controller ports are working and correctly configured in the controller and the SAN zoning. Failure to do this can result in errors that are not reported.
 - b. Although it might appear that the **detectmdisk** command has completed, extra time might be required for it to run. The **detectmdisk** is asynchronous and returns a prompt while the command continues to run in the background. You can use the **lsdiscoverystatus** command to view the discovery status.
2. When the detection is complete, issue the **svcinfolsmdiskcandidate** CLI command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group (storage pools).
 3. Issue the **svcinfolsmdisk** CLI command to view all of the MDisks.

You have now seen that the back-end controllers and switches have been set up correctly and that the SAN Volume Controller cluster recognizes the storage that is presented by the back-end controller.

The following example describes a scenario where a single back-end controller is presenting eight SCSI LUs to the SAN Volume Controller cluster:

1. Issue `svctask detectmdisk`.
2. Issue `svcinfolsmdiskcandidate`.

The following output is displayed:

```
id
0
1
2
3
4
5
6
7
```

3. Issue `svcinfolsmdisk -delim :`

The following output is displayed:

```

|   svcinfo lsmdisk -delim :
|   id:name:status:mode:mdisk_grp_id:mdisk_grp_name:capacity:ctrl_LUN_#:controller_name:UID:tier
|   0:mdisk0:online:unmanaged:::68.4GB:0000000000000000:controller0:
|   20000004cf2422aa000000000000000000000000000000000000000000000000:
|   1:mdisk1:online:unmanaged:::68.4GB:0000000000000000:controller1:
|   20000004cf1fd19d000000000000000000000000000000000000000000000000:
|   2:mdisk2:online:unmanaged:::68.4GB:0000000000000000:controller2:
|   20000004cf242531000000000000000000000000000000000000000000000000:

```

Creating MDisk groups (storage pools) using the CLI

You can use the command-line interface (CLI) to create an MDisk group (storage pool).

Attention: If you add an MDisk to an MDisk group (storage pool) as an MDisk, any data on the MDisk is lost. If you want to keep the data on an MDisk (for example, because you want to import storage that was previously not managed by SAN Volume Controller), you must create image mode VDIsks (volumes) instead.

Assume that the cluster has been set up and that a back-end controller has been configured to present new storage to the SAN Volume Controller.

If you are using a SAN Volume Controller solid-state drive (SSD) managed disk, ensure that you are familiar with the SSD configuration rules.

If you intend to keep the VDisk (volume) allocation within one storage system, ensure that all MDisks in the MDisk group (storage pool) are presented by the same storage system.

Ensure that all MDisks that are allocated to a single MDisk group (storage pool) are of the same RAID type. If the MDisk group (storage pool) has more than one tier of storage, ensure that all MDisks in the same tier are of the same RAID type. When using Easy Tier™, all of the MDisks in an MDisk group (storage pool) in the same tier should be similar and have similar performance characteristics. If you do not use Easy Tier, the MDisk group (storage pool) should contain only one tier of storage, and all of the MDisks in the MDisk group (storage pool) should be similar and have similar performance characteristics.

Consider the following factors as you decide how many MDisk groups (storage pools) to create:

- A VDisk (volume) can only be created using the storage from one MDisk group (storage pool). Therefore, if you create small MDisk groups (storage pools), you might lose the benefits that are provided by virtualization, namely more efficient management of free space and a more evenly distributed workload for better performance.
- If any MDisk in an MDisk group (storage pool) goes offline, all the VDIsks (volumes) in the MDisk group (storage pool) go offline. Therefore you might want to consider using different MDisk groups (storage pools) for different back-end controllers or for different applications.
- If you anticipate regularly adding and removing back-end controllers or storage, this task is made simpler by grouping all the MDisks that are presented by a back-end controller into one MDisk group (storage pool).
- All the MDisks in an MDisk group (storage pool) should have similar levels of performance or reliability, or both. If an MDisk group (storage pool) contains MDisks with different levels of performance, the performance of the VDIsks (volumes) in this group is limited by the performance of the slowest MDisk. If an MDisk group (storage pool) contains MDisks with different levels of reliability, the reliability of the VDIsks (volumes) in this group is that of the least reliable MDisk in the group.

Note: When you create a storage pool with a new solid-state drive (SSD), the new SSD is automatically formatted and set to a block size of 512 bytes.

Even with the best planning, circumstances can change and you must reconfigure your MDisk groups (storage pools) after they have been created. The data migration facilities that are provided by the SAN Volume Controller enable you to move data without disrupting I/O.

Choosing an MDisk group (storage pool) extent size

Consider the following factors as you decide the extent size of each new MDisk group (storage pool):

- You must specify the extent size when you create a new MDisk group (storage pool).
- You cannot change the extent size later; it must remain constant throughout the lifetime of the MDisk group (storage pool).
- MDisk groups (storage pools) can have different extent sizes; however, this places restrictions on the use of data migration.
- The choice of extent size affects the maximum size of a VDisk (volume) in the MDisk group (storage pool).

Table 4 compares the maximum VDisk (volume) capacity for each extent size. The maximum is different for thin-provisioned volumes. Because the SAN Volume Controller allocates a whole number of extents to each VDisk (volume) that is created, using a larger extent size might increase the amount of storage that is wasted at the end of each VDisk (volume). Larger extent sizes also reduces the ability of the SAN Volume Controller to distribute sequential I/O workloads across many MDisks and therefore can reduce the performance benefits of virtualization.

Table 4. Maximum VDisk (volume) capacity by extent size

Extent size (MB)	Maximum VDisk (volume) capacity in GB (not thin-provisioned volumes)	Maximum VDisk (volume) capacity in GB (thin-provisioned volumes)
16	2048 (2 TB)	2000
32	4096 (4 TB)	4000
64	8192 (8 TB)	8000
128	16,384 (16 TB)	16,000
256	32,768 (32 TB)	32,000
512	65,536 (64 TB)	65,000
1024	131,072 (128 TB)	130,000
2048	262,144 (256 TB)	260,000
4096	262,144 (256 TB)	262,144
8192	262,144 (256 TB)	262,144

Important: You can specify different extent sizes for different MDisk groups (storage pools); however, you cannot migrate VDIs (volumes) between MDisk groups (storage pools) with different extent sizes. If possible, create all your MDisk groups (storage pools) with the same extent size.

Perform the following steps to create an MDisk group (storage pool):

Issue the `svctask mkmdiskgrp` CLI command to create an MDisk group (storage pool).

The following is an example of the CLI command you can issue to create an MDisk group (storage pool):

```
svctask mkmdiskgrp -name maindiskgroup -ext 32
-mdisk mdisk0:mdisk1:mdisk2:mdisk3
```

| where *maindiskgroup* is the name of the MDisk group (storage pool) that you want to create, 32 MB is the size of the extent you want to use, and *mdsk0*, *mdsk1*, *mdsk2*, *mdsk3* are the names of the four MDisks that you want to add to the group.

| You created and added MDisks to an MDisk group (storage pool).

| The following example provides a scenario where you want to create an MDisk group (storage pool), but you do not have any MDisks available to add to the group. You plan to add the MDisks at a later time.

| You use the **svctask mkmdiskgrp** CLI command to create the MDisk group (storage pool) *bkpmdiskgroup* and later used the **svctask addmdisk** CLI command to add *mdsk4*, *mdsk5*, *mdsk6*, *mdsk7* to the MDisk group (storage pool).

1. Issue `svctask mkmdiskgrp -name bkpmdiskgroup -ext 32`

| where *bkpmdiskgroup* is the name of the MDisk group (storage pool) that you want to create and 32 MB is the size of the extent that you want to use.

| 2. You find four MDisks that you want to add to the MDisk group (storage pool).

3. Issue `svctask addmdisk -mdisk msk4:msk5:msk6:msk7 bkpmdiskgroup`

| where *mdsk4*, *mdsk5*, *mdsk6*, *mdsk7* are the names of the MDisks that you want to add to the MDisk group (storage pool) and *bkpmdiskgroup* is the name of the MDisk group (storage pool) for which you want to add MDisks.

| Adding MDisks to MDisk groups (storage pools) using the CLI

| You can use the command-line interface (CLI) to add managed disks (MDisks) to MDisk groups (storage pools).

| The MDisks must be in unmanaged mode. Disks that already belong to an MDisk group (storage pool) cannot be added to another MDisk group (storage pool) until they have been deleted from their current MDisk group (storage pool). You can delete an MDisk from an MDisk group (storage pool) under the following circumstances:

- | • If the MDisk does not contain any extents in use by a virtual disk (VDisk) (volume)
- | • If you can first migrate the extents in use onto other free extents within the group

| **Important:** Do not add an MDisk using this procedure if you are mapping the MDisk to an image mode VDisk. Adding an MDisk to an MDisk group (storage pool) allows the SAN Volume Controller to write new data to the MDisk; therefore, any existing data on the MDisk is lost. If you want to create an image mode VDisk, use the **svctask mkvdisk** command instead of **svctask addmdisk**.

| If you are using a SAN Volume Controller solid-state drive (SSD) managed disk, ensure that you are familiar with the SSD configuration rules.

| When you are adding MDisks to an MDisk group (storage pool) using the **svctask addmdisk** command or when you are creating an MDisk group (storage pool) using the **svctask mkmdiskgrp -mdisk** command, the SAN Volume Controller performs tests on the MDisks in the list before the MDisks are allowed to become part of an MDisk group (storage pool). These tests include checks of the MDisk identity, capacity, status and the ability to perform both read and write operations. If these tests fail or exceed the time allowed, the MDisks are not added to the group. However, with the **svctask mkmdiskgrp -mdisk** command, the MDisk group (storage pool) is still created even if the tests fail, but it does not contain any MDisks. If tests fail, confirm that the MDisks are in the correct state and that they have been correctly discovered.

| The following events contribute to an MDisk test failure:

- | • The MDisk is not visible to all SAN Volume Controller nodes in the cluster.
- | • The MDisk identity has changed from a previous discovery operation.

- The MDisk cannot perform read or write operations.
- The status of the MDisk can be either degraded paths, degraded ports, excluded, or offline.
- The MDisk does not exist.

The following events contribute to an MDisk test timeout:

- The disk controller system on which the MDisk resides is failing.
- A SAN fabric or cable fault condition exists that is preventing reliable communication with the MDisk.

Note: The first time that you add a new solid-state drive (SSD) to a storage pool, the SSD is automatically formatted and set to a block size of 512 bytes.

1 Perform the following steps to add MDisks to MDisk groups (storage pools):

1. Issue the **svcinfolsmdiskgrp** CLI command to list the existing MDisk groups (storage pools).

The following is an example of the CLI command you can issue to list the existing MDisk groups (storage pools):

```
svcinfolsmdiskgrp -delim :
```

The following is an example of the output that is displayed:

```
id:name:status:mdisk_count:vdisk_count:
capacity:extent_size:free_capacity:virtual_capacity:
used_capacity:real_capacity:overallocation:warning
0:mdiskgrp0:online:3:4:33.3GB:16:32.8GB:64.00MB:64.00MB:64.00MB:0:0
1:mdiskgrp1:online:2:1:26.5GB:16:26.2GB:16.00MB:16.00MB:16.00MB:0:0
2:mdiskgrp2:online:2:0:33.4GB:16:33.4GB:0.00MB:0.00MB:0.00MB:0:0
```

2. Issue the **svctask addmdisk** CLI command to add MDisks to the MDisk group (storage pool).

The following is an example of the CLI command you can issue to add MDisks to an MDisk group (storage pool):

```
svctask addmdisk -mdisk mdisk4:mdisk5:mdisk6:mdisk7 bkpmdiskgroup
```

Where *mdisk4:mdisk5:mdisk6:mdisk7* are the names of the MDisks that you want to add to the MDisk group (storage pool) and *bkpmdiskgroup* is the name of the MDisk group (storage pool) for which you want to add the MDisks.

Setting a quorum disk using the CLI

You can set a managed disk (MDisk) as a quorum disk by using the command-line interface (CLI). With Storwize V7000 you can also set a drive as a quorum drive.

To set an MDisk as the active quorum disk, use the **chquorum** command. With Storwize V7000, only assign quorum to drives in the control enclosure or to external MDisks.

- When possible, distribute the quorum candidate disks so that each MDisk is provided by a different storage system. For a list of storage systems that support quorum disks, refer to the Supported Hardware List located at the Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145.
- Ensure that the status of the MDisk that is being set as a quorum disk is online before setting it as the quorum disk. To check the status of the MDisk that is being set as the quorum disk, use the **svcinfolsqorum** command.
- Wait at least 20 seconds between issuing a consecutive **chquorum** command. This delay allows time for the first update to the quorum disk to complete before setting another quorum disk.

Note: Quorum functionality is not supported for solid-state drive (SSD) MDisks.

To set an MDisk as a quorum disk, follow these steps:

1. Issue the following CLI command to specify the MDisk or drive to be a specified quorum ID. The cluster automatically assigns quorum indexes.


```
svctask chquorum -mdisk mdisk_id_or_name quorum_id
svctask chquorum -drive drive_id quorum_id
```

 where *mdisk_id_or_name* or *drive_id* is the ID or the name of the MDisk or drive to assign as a quorum candidate disk or drive.
2. Issue the following command to make the specified MDisk or drive the active quorum disk or drive. To identify the quorum disk that is active, use the `svcinfolsqorum` command.


```
svctask chquorum -active mdisk_id_or_name quorum_id
svctask chquorum -active drive_id quorum_id
```

 where *mdisk_id_or_name* or *drive_id* is the ID or the name of the MDisk or drive to assign as a quorum candidate disk or drive. The **active** parameter makes the MDisk that was specified in the previous example the active quorum disk.

Modifying the amount of available memory for Copy Services and VDisk Mirroring features using the CLI

You can use the command-line interface (CLI) to modify the amount of memory that is available for the VDisk (Volume) Mirroring feature and the FlashCopy, Metro Mirror, or Global Mirror Copy Services features.

Table 5 provides an example of the amount of memory that is required for VDisk Mirroring and each Copy Service feature.

Table 5. Memory required for VDisk Mirroring and Copy Services

Feature	Grain size	1 MB of memory provides the following VDisk capacity for the specified I/O group
Metro Mirror or Global Mirror	256 KB	2 TB of total Metro Mirror and Global Mirror VDisk capacity
FlashCopy	256 KB	2 TB of total FlashCopy source VDisk capacity
FlashCopy	64 KB	512 GB of total FlashCopy source VDisk capacity
Incremental FlashCopy	256 KB	1 TB of total incremental FlashCopy source VDisk capacity
Incremental FlashCopy	64 KB	256 GB of total incremental FlashCopy source VDisk capacity
VDisk Mirroring	256 KB	2 TB of mirrored VDisk capacity
Notes:		
<ol style="list-style-type: none"> 1. For multiple FlashCopy targets, you must consider the number of mappings. For example, for a mapping with a grain size of 256 KB, 8 KB of memory allows one mapping between a 16 GB source VDisk and a 16 GB target VDisk. Alternatively, for a mapping with a 256 KB grain size, 8 KB of memory allows two mappings between one 8 GB source VDisk and two 8 GB target VDIs. 2. When creating a FlashCopy mapping, if you specify an I/O group other than the I/O group of the source VDisk, the memory accounting goes towards the specified I/O group, not towards the I/O group of the source VDisk. 3. For VDisk Mirroring, the full 512 MB of memory space provides 1 PB of total VDisk Mirroring capacity. 4. In this table, <i>capacity</i> refers to the virtual capacity of the VDisk. For thin-provisioned volumes with different virtual capacities and real capacities, the virtual capacity is used for memory accounting. 		

Table 6 on page 24 provides an example of RAID level comparisons with their bitmap memory cost, where *MS* is the size of the member drives and *MC* is the number of member drives.

Table 6. RAID level comparisons

Level	Member count	Approximate capacity	Redundancy	Approximate bitmap memory cost
RAID-0	1-8	MC * MS	None	(1 MB per 2 TB of MS) * MC
RAID-1	2	MS	1	(1 MB per 2 TB of MS) * (MC/2)
RAID-5	3-16	(MC-1) * MS	1	1 MB per 2 TB of MS with a strip size of 256 KB; double with strip size of 128 KB.
RAID-6	5-16	less than (MC-2 * MS)	2	
RAID-10	2-16 (evens)	MC/2 * MS	1	(1 MB per 2 TB of MS) * (MC/2)

Note: There is a margin of error on the approximate bitmap memory cost of approximately 15%. For example, the cost for a 256 KB RAID-5 is ~1.15 MB for the first 2 TB of MS.

To modify and verify the amount of memory that is available, perform the following steps:

1. Issue the following command to modify the amount of memory that is available for VDisk Mirroring or a Copy Service feature:

```
svctask chiogrp -feature flash|remote|mirror -size memory_size io_group_id | io_group_name
```

where *flash|remote|mirror* is the feature that you want to modify, *memory_size* is the amount of memory that you want to be available, and *io_group_id* | *io_group_name* is the ID or name of the I/O group for which you want to modify the amount of available memory.

2. Issue the following command to verify that the amount of memory has been modified:

```
svcinfo lsiogrp object_id | object_name
```

where *object_id* | *object_name* is the ID or name of the I/O group for which you have modified the amount of available memory.

The following information is an example of the output that is displayed.

```
id 0
name io_grp 0
node_count 2
vdisk_count 28
host_count 2
flash_copy_total_memory 20.0MB
flash_copy_free_memory 20.0MB
remote_copy_total_memory 20.0MB
remote_copy_free_memory 20.0MB
mirroring_total_memory 10.0MB
mirroring_free_memory 10.0MB
raid_total_memory 20.0MB
raid_free_memory 19.2MB
maintenance no <---
```

Creating VDisks (volumes) using the CLI

You can use the command-line interface (CLI) to create a virtual disk (VDisk) (volume).

If the VDisk (volume) that you are creating maps to a solid-state drive (SSD), the data that is stored on the VDisk (volume) is not protected against SSD failures or node failures. To avoid data loss, add a VDisk (volume) copy that maps to an SSD on another node.

This task assumes that the cluster has been set up and that you have created managed disk (MDisk) groups (storage pools). You can establish an empty MDisk group (storage pool) to hold the MDisks that are used for image mode VDisks (volumes).

| **Note:** If you want to keep the data on an MDisk, create image mode VDisks (volumes). This task describes how to create a VDisk (volume) with striped virtualization.

| Perform the following steps to create VDisks (volumes):

1. Issue the **svcinfolsmdiskgrp** CLI command to list the available MDisk groups (storage pools) and the amount of free storage in each group.

| The following is an example of the CLI command you can issue to list MDisk groups (storage pools):

svcinfolsmdiskgrp -delim :

The following is an example of the output that is displayed:

```
id:name:status:mdisk_count:vdisk_count:capacity:extent_size:free_capacity:virtual_capacity:
used_capacity:real_capacity:overallocation:warning:easy_tier:easy_tier_status
0:mdiskgrp0:degraded:4:0:34.2GB:16:34.2GB:0:0:0:0:auto:inactive
1:mdiskgrp1:online:4:6:200GB:16:100GB:400GB:75GB:100GB:200:80:on:active
```

2. Decide which MDisk group (storage pool) you want to provide the storage for the VDisk (volume).
3. Issue the **svcinfolsiogrp** CLI command to show the I/O groups and the number of VDisks (volumes) assigned to each I/O group.

| **Note:** It is normal for clusters with more than one I/O group to have MDisk groups (storage pools) that have VDisks (volumes) in different I/O groups. You can use FlashCopy to make copies of VDisks (volumes) regardless of whether the source and target VDisk (volume) are in the same I/O group. If you plan to use intracluster Metro Mirror or Global Mirror, both the master and auxiliary VDisk (volume) must be in the same I/O group.

The following is an example of the CLI command you can issue to list I/O groups:

svcinfolsiogrp -delim :

The following is an example of the output that is displayed:

```
id:name:node_count:vdisk_count:host_count
0:io_grp0:2:0:2
1:io_grp1:2:0:1
2:io_grp2:0:0:0
3:io_grp3:0:0:0
4:recovery_io_grp:0:0:0
```

4. Decide which I/O group you want to assign the VDisk (volume) to. This determines which SAN Volume Controller nodes in the cluster process the I/O requests from the host systems. If you have more than one I/O group, make sure you distribute the VDisks (volumes) between the I/O groups so that the I/O workload is shared evenly between all SAN Volume Controller nodes.
5. Issue the **svctask mkvdisk** CLI command to create a VDisk (volume).
The rate at which the VDisk (volume) copies will resynchronize after loss of synchronization can be specified using the **syncrate** parameter. The following table defines the rates:

Table 7. VDisk (volume) copy resynchronization rates

Syncrate value	Data copied per second
1-10	128 KB
11-20	256 KB
21-30	512 KB
31-40	1 MB
41-50	2 MB
51-60	4 MB
61-70	8 MB
71-80	16 MB

Table 7. VDisk (volume) copy resynchronization rates (continued)

Syncrate value	Data copied per second
81-90	32 MB
91-100	64 MB

The default setting is 50. The synchronization rate must be set such that the VDisk (volume) copies will resynchronize quickly after loss of synchronization.

When a mirrored VDisk (volume) uses disk extents on a solid-state drive (SSD) that is located on a SAN Volume Controller node, synchronization is lost if one of the nodes goes offline either during a concurrent code upgrade or because of maintenance. During code upgrade, the synchronization must be restored within 30 minutes or the upgrade stalls. Unlike VDisk (volume) copies from external storage systems, during the period that the SSD VDisk (volume) copies are not synchronized, access to the VDisk (volume) depends on the single node containing the SSD storage associated with the synchronized VDisk (volume) copy. The default synchronization rate is typically too low for SSD VDisk (volume) mirrors. Instead, set the synchronization rate to 80 or above.

The following is an example of the CLI command that you can issue to create a VDisk (volume) with two copies using the I/O group and MDisk group (storage pool) name and specifying the synchronization rate:

```
svctask mkvdisk -iogrp io_grp1 -mdiskgrp grpa:grpb -size500 -vtype striped
-copies 2 -syncrate 90
```

where *io_grp1* is the name of the I/O group that you want the VDisk (volume) to use, *grpa* is the name of the MDisk group (storage pool) for the primary copy of the VDisk (volume) and *grpb* is the name of the MDisk group (storage pool) for the second copy of the VDisk (volume), and 2 is the number of VDisk (volume) copies and the synchronization rate is 90 which is equivalent to 32MB per second.

The following is an example of the CLI command you can issue to create a VDisk (volume) using the I/O group ID and MDisk group (storage pool) ID:

```
svctask mkvdisk -name mainvdisk1 -iogrp 0
-mdiskgrp 0 -vtype striped -size 256 -unit gb
```

where *mainvdisk1* is the name that you want to call the VDisk (volume), 0 is the ID of the I/O group that want the VDisk (volume) to use, 0 is the ID of the MDisk group (storage pool) that you want the VDisk (volume) to use, and 256 is the capacity of the VDisk (volume).

The following is an example of the CLI command that you can issue to create a VDisk (volume) using the I/O group and MDisk group (storage pool) name:

```
svctask mkvdisk -name bkpvdisk1 -iogrp io_grp1
-mdiskgrp bkpmdiskgroup -vtype striped -size 256 -unit gb
```

where *bkpvdisk1* is the name that you want to call the VDisk (volume), *io_grp1* is the name of the I/O group that want the VDisk (volume) to use, *bkpmdiskgroup* is the name of the MDisk group (storage pool) that you want the VDisk (volume) to use, and 256 is the capacity of the VDisk (volume).

The following is an example of the CLI command that you can issue to create a space-efficient VDisk (volume) using the I/O group and MDisk group (storage pool) name:

```
svctask mkvdisk -iogrp io_grp1 -mdiskgrp bkpmdiskgroup -vtype striped
-size 10 unit gb -rsize 20% -autoexpand -grainsize 32
```

where *io_grp1* is the name of the I/O group that you want the VDisk (volume) to use and 20% is how much real storage to allocate to the VDisk (volume), as a proportion of its virtual size. In this example, the size is 10 GB so that 2 GB will be allocated.

The following is an example of the CLI command that you can issue to create a VDisk (volume) with two copies using the I/O group and MDisk group (storage pool) name:

```
svctask mkvdisk -iogrp io_grp1 -mdiskgrp grpa:grpb
-size 500 -vtype striped -copies 2
```

| where *io_grp1* is the name of the I/O group that you want the VDisk (volume) to use, *grpa* is the
| name of the MDisk group (storage pool) for the primary copy of the VDisk (volume) and *grp b* is the
| name of the MDisk group (storage pool) for the second copy of the VDisk (volume), and 2 is the
| number of VDisk (volume) copies.

| **Note:** If you want to create two VDisk (volume) copies of different types, create the first copy using
| the **mkvdisk** command and then add the second copy using the **addvdiskcopy** command.

| 6. Issue the **svcinfo lsvdisk** CLI command to list all the VDIsks (volumes) that have been created.

Adding a copy to a VDisk (volume) using the CLI

| You can use the command-line interface (CLI) to add a mirrored copy to a virtual disk (VDisk). Each
| VDisk can have a maximum of two copies.

The **addvdiskcopy** command adds a copy to an existing VDisk, which changes a nonmirrored VDisk into a mirrored VDisk.

| Creating mirrored copies of a VDisk allows the VDisk to remain accessible even when a managed disk
| (MDisk) that the VDisk depends on becomes unavailable. You can create copies of a VDisk either from
| different MDisk groups or by creating an image mode copy of the VDisk. Copies allow for availability of
| data; however, they are not separate objects. You can only create or change mirrored copies from the
| VDisk.

| In addition, you can use VDisk mirroring as an alternative method of migrating VDIsks between MDisk
| groups. For example, if you have a nonmirrored VDisk in one MDisk group and want to migrate that
| VDisk to a second MDisk group, you can add a new copy of the VDisk by specifying the second MDisk
| group for that VDisk copy. After the copies have synchronized, you can delete the copy in the first MDisk
| group. The VDisk is migrated to the second MDisk group while remaining online during the migration.

| This alternative method of migrating VDIsks has the following advantages:

- | • Access to the VDisk data is not lost if the second MDisk group goes offline during the migration.
- | • The speed of the migration can be adjusted, using the VDisk synchronization rate, and the migration can be paused.
- | • The migration can be ended by deleting the VDisk copy in the second MDisk group before migration completes.
- | • The MDisk groups can have different extent sizes.

| This alternative method has the following limitations:

- | • You cannot use this method for VDIsks that are already mirrored.
- | • There are more manual steps that are associated with this method.
- | • Write I/O performance is slightly affected during the migration, because the mirrored copies must be kept synchronized.

Use the **-copies** parameter to specify the number of copies to add to the VDisk; this is currently limited to the default value of 1 copy. Use the **-mdiskgrp** parameter to specify the managed disk group that will provide storage for the copy; the **svcinfo lsmdiskgrp** CLI command lists the available managed disk groups and the amount of available storage in each group.

For image copies, you must specify the virtualization type using the **-vtype** parameter and an MDisk that is in unmanaged mode using the **-mdisk** parameter. This MDisk must be in the unmanaged mode. The **-vtype** parameter is optional for sequential (seq) and striped VDIsks. The default virtualization type is **striped**.

Use the **syncrate** parameter to specify the rate at which the VDisk copies will resynchronize after loss of synchronization. The topic that describes creating VDIsks (volumes) using the CLI describes this parameter.

Issue the **addvdiskcopy** CLI command to add a mirrored copy to a VDisk:

```
svctask addvdiskcopy -mdiskgrp 0 vdisk8
```

where *0* is the name of the managed disk group and *vdisk8* is the VDisk to which the copy will be added.

The command returns the IDs of the newly created VDisk copies.

Deleting a copy from a VDisk (volume) using the CLI

You can use the command-line interface (CLI) to delete a mirrored copy from a virtual disk (VDisk) (volume).

If you are using solid-state drives (SSDs) that are inside a SAN Volume Controller node, always use VDisk (volume) mirroring with these SSDs. Data stored on SSDs inside SAN Volume Controller nodes is not protected against SSD failures or node failures. Therefore, if you are deleting a VDisk (volume) copy that uses a SSD, ensure that the data that is stored on the copy is protected with another VDisk (volume) copy.

The **rmvdiskcopy** CLI command deletes the specified copy from the specified VDisk (volume). The command fails if all other copies of the VDisk (volume) are not synchronized; in this case, you must specify the **-force** parameter, delete the VDisk (volume), or wait until the copies are synchronized. You must specify the **vdisk_name** | **vdisk_id** parameter last on the command line.

Issue the **rmvdiskcopy** CLI command to delete a mirrored copy from a VDisk (volume):

```
svctask rmvdiskcopy -copy 1 vdisk8
```

where *1* is the ID of the copy to delete and *vdisk8* is the virtual disk to delete the copy from.

The command does not return any output.

Configuring host objects using the CLI

You can use command-line interface (CLI) to create host objects.

If you are configuring a host object on a fibre-channel attached host, ensure that you have completed all zone and switch configuration. Also test the configuration to ensure that zoning was created correctly.

If you are configuring a host object on the cluster that uses iSCSI connections, ensure that you have completed the necessary host-system configurations and have configured the cluster for iSCSI connections.

At least one WWPN or iSCSI name must be specified.

Perform the following steps to create host objects:

1. Issue the **svctask mkhost** CLI command to create a logical host object for a fibre-channel attached host. Assign your worldwide port name (WWPN) for the host bus adapters (HBAs) in the hosts.

The following is an example of the CLI command that you can issue to create a fibre-channel attached host:

```
svctask mkhost -name new_name -hbawwpn wwpn_list
```

where *new_name* is the name of the host and *wwpn_list* is the WWPN of the HBA.

2. To create an iSCSI-attached host, issue the following CLI command:

```
svctask mkhost -iscsiname iscsi_name_list
```

where *iscsi_name_list* specifies one or more iSCSI qualified names (IQNs) of this host. Up to 16 names can be specified, provided that the command-line limit is not reached. Each name should comply with the iSCSI standard, RFD 3720.

3. To add ports to a fibre-channel attached host, issue the **svctask addhostport** CLI command.

For example, issue the following CLI command:

```
svctask addhostport -hbawwpn wwpn_list new_name
```

This command adds another HBA WWPN *wwpn_list* to the host that was created in step 1 on page 28.

4. To add ports to an iSCSI-attached host, issue the **svctask addhostport** CLI command.

For example, issue the following CLI command:

```
svctask addhostport -iscsiname iscsi_name_list new_name
```

where *iscsi_name_list* Specifies the comma-separated list of IQNs to add to the host. This command adds an IQN to the host that was created in step 2.

5. To set the Challenge Handshake Authentication Protocol (CHAP) secret that is used to authenticate the host for iSCSI I/O, issue the **svctask chhost** CLI command. This secret is shared between the host and the cluster. For example, issue the following CLI command:

```
svctask chhost -chapsecret chap_secret
```

where *chap_secret* is the CHAP secret that is used to authenticate the host for iSCSI I/O. To list the CHAP secret for each host, use the **svcinfo lsiscsiauth** command. To clear any previously set CHAP secret for a host, use the **svctask chhost -nochapsecret** command.

| After you have created the host object on the cluster, you can map volumes to a host.

If you are unable to discover the disk on the host system or if there are fewer paths available for each disk than expected, test the connectivity between your host system and the cluster. Depending on the connection type to the host, these steps might be different. For iSCSI-attached hosts, test your connectivity between the host and SAN Volume Controller ports by pinging SAN Volume Controller from the host. Ensure that the firewall and router settings are configured correctly and validate that the values for the subnet mask and gateway are specified correctly for the SAN Volume Controller host configuration.

For fibre-channel attached hosts, ensure that the active switch configuration includes the host zone and check the host-port link status. To verify end-to-end connectivity, you can use the **svcinfo lsfabric** CLI command or the View Fabric panel under Service and Maintenance container in the management GUI.

Creating VDisk-to-host mappings using the CLI

You can use the command-line interface (CLI) to create virtual disk (VDisk)-to-host mappings.

- | To create VDisk-to-host mappings (host mappings), follow these steps:
- | 1. Issue the **svctask mkvdiskhostmap** CLI command to create VDisk-to-host mappings.
The following is an example of the CLI command you can issue to create VDisk-to-host mappings:

```
svctask mkvdiskhostmap -host demohost1 mainvdisk1
```


Where *demohost1* is the name of the host and *mainvdisk1* is the name of the VDisk.
 - | 2. After you have mapped VDIs to hosts, discover the disks on the host system. This step requires that
| you access the host system and use the host-system utilities to discover the new disks that are made
| available by the SAN Volume Controller. You also have the option of creating a file system for those
| new disks. Consult your host-system documentation for more information on completing this task.

Creating FlashCopy mappings using the CLI

You can use the command-line interface (CLI) to create FlashCopy mappings.

- | A FlashCopy mapping specifies the source and target virtual disk (VDisk) (volume). Source VDIs
| (volumes) and target VDIs (volumes) must meet the following requirements:
- They must be the same size.
 - They must be managed by the same cluster.
- | A VDisk (volume) can be the source in up to 256 mappings. A mapping is started at the point in time
| when the copy is required.

Perform the following steps to create FlashCopy mappings:

- | 1. The source and target VDisk (volume) must be the exact same size. Issue the **svcinfo lsvdisk -bytes**
| CLI command to find the size (capacity) of the VDisk (volume) in bytes.
- | 2. Issue the **svctask mkfcmap** CLI command to create a FlashCopy mapping.

The following CLI command example creates a FlashCopy mapping and sets the copy rate:

```
svctask mkfcmap -source mainvdisk1 -target bkpvdisk1  
-name main1copy -copyrate 75
```

- | Where *mainvdisk1* is the name of the source VDisk (volume), *bkpvdisk1* is the name of the VDisk
| (volume) that you want to make the target VDisk (volume), *main1copy* is the name that you want to
| call the FlashCopy mapping, and 75 is the copy rate.

The following is an example of the CLI command you can issue to create FlashCopy mappings
without the copy rate parameter:

```
svctask mkfcmap -source mainvdisk2 -target bkpvdisk2  
-name main2copy
```

- | Where *mainvdisk2* is the name of the source VDisk (volume), *bkpvdisk2* is the name of the VDisk
| (volume) that you want to make the target VDisk (volume), *main2copy* is the name that you want to
| call the FlashCopy mapping.

Note: The default copy rate of 50 is used if you do not specify a copy rate.

- | If the specified source and target VDIs (volume) are also the target and source VDIs (volumes) of
| an existing mapping, the mapping that is being created and the existing mapping become partners. If
| one mapping is created as incremental, its partner is automatically incremental. A mapping can have
| only one partner.
- | 3. Issue the **svcinfo lsfcmap** CLI command to check the attributes of the FlashCopy mappings that have
| been created:

The following is an example of a CLI command that you can issue to view the attributes of the FlashCopy mappings:

```
svcinfolsfcmapp -delim :
```

Where **-delim** species the delimiter. The following is an example of the output that is displayed:

```
id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:target_vdisk_name:
group_id:group_name:status:progress:copy_rate:clean_progress:incremental
0:main1copy:77:vdisk77:78:vdisk78:::idle_or_copied:0:75:100:off
1:main2copy:79:vdisk79:80:vdisk80:::idle_or_copied:0:50:100:off
```

Preparing and starting a FlashCopy mapping using the CLI

Before you start the FlashCopy process using the command-line interface (CLI), you must prepare a FlashCopy mapping.

Starting a FlashCopy mapping creates a point-in-time copy of the data on the source virtual disk (VDisk) and writes it to the target VDisk (volume) for the mapping.

Perform the following steps to prepare and start a FlashCopy mapping:

1. Issue the **svctask prestartfcmap** CLI command to prepare the FlashCopy mapping.

To run the following command, the FlashCopy mapping cannot belong to a consistency group.

```
svctask prestartfcmap -restore main1copy
```

Where *main1copy* is the name of the FlashCopy mapping.

This command specifies the optional **restore** parameter, which forces the mapping to be prepared even if the target VDisk is being used as a source in another active FlashCopy mapping.

The mapping enters the preparing state and moves to the prepared state when it is ready.

2. Issue the **svcinfolsfcmapp** CLI command to check the state of the mapping.

The following is an example of the output that is displayed:

```
svcinfolsfcmapp -delim :
id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:
target_vdisk_name:group_id:group_name:status:progress:copy_rate
0:main1copy:0:mainvdisk1:1:bkpvdisk1:::prepared:0:50
```

3. Issue the **svctask startfcmap** CLI command to start the FlashCopy mapping.

The following is an example of the CLI command you can issue to start the FlashCopy mapping:

```
svctask startfcmap -restore main1copy
```

Where *main1copy* is the name of the FlashCopy mapping.

This command specifies the optional **restore** parameter, which forces the mapping to be started even if the target VDisk is being used as a source in another active FlashCopy mapping.

4. Issue the **svcinfolsfcmappprogress** CLI command with the FlashCopy mapping name or ID to check the progress of the mapping.

The following is an example of the output that is displayed; the FlashCopy mapping ID 0 is 47% completed.

```
svcinfolsfcmappprogress -delim :
id:progress
0:47
```

You have created a point-in-time copy of the data on a source VDisk and written that data to a target VDisk. The data on the target VDisk is only recognized by the hosts that are mapped to it.

Stopping FlashCopy mappings using the CLI

You can use the command-line interface (CLI) to stop a FlashCopy mapping.

Follow these steps to stop a single stand-alone FlashCopy mapping.

1. To stop a FlashCopy mapping, issue the following **svctask stopfcmap** command:

```
svctask stopfcmap fc_map_id or fc_map_name
```

where *fc_map_id* or *fc_map_name* is the ID or name of the mapping to stop.

2. To stop immediately all processing associated with the mapping and break the dependency on the source VDisk (volume) of any mappings that are also dependent on the target disk, issue the following command:

```
svctask stopfcmap -force -split fc_map_id or fc_map_name
```

When you use the **force** parameter, all FlashCopy mappings that depend on this mapping (as listed by the **lsfcmapdependentmaps** command) are also stopped. The **split** parameter can be specified only when stopping a map that has a progress of 100 as shown by the **svcinfo lsfcmap** command. The **split** parameter removes the dependency of any other mappings on the source VDisk. It might be used prior to starting another FlashCopy mapping whose target disk is the source disk of the mapping being stopped. After the mapping is stopped with the **split** option, you can start the other mapping without the **restore** option.

Deleting a FlashCopy mapping using the CLI

You can use the command-line interface (CLI) to delete a FlashCopy mapping.

The **svctask rmfcmap** CLI command deletes an existing mapping if the mapping is in the `idle_or_copied` or `stopped` state. If it is in the `stopped` state, the **force** parameter is required to specify that the target VDisk (volume) is brought online. If the mapping is in any other state, you must stop the mapping before you can delete it.

If deleting the mapping splits the tree that contains the mapping, none of the mappings in either resulting tree can depend on any mapping in the other tree. To display a list of dependent FlashCopy mappings, use the **svcinfo lsfcmapdependentmaps** command.

1. To delete an existing mapping, issue the **rmfcmap** CLI command:

```
svctask rmfcmap fc_map_id or fc_map_name
```

where *fc_map_id* or *fc_map_name* is the ID or name of the mapping to delete.

2. To delete an existing mapping and bring the target VDisk online, issue the following command:

```
svctask rmfcmap -force fc_map_id or fc_map_name
```

where *fc_map_id* or *fc_map_name* is the ID or name of the mapping to delete.

The command does not return any output.

Creating a FlashCopy consistency group and adding mappings using the CLI

You can use the command-line interface (CLI) to create and add mappings to a FlashCopy consistency group.

1. If you have created several FlashCopy mappings for a group of virtual disks (volumes) that contain elements of data for the same application, it can be convenient to assign these mappings to a single FlashCopy consistency group. You can then issue a single prepare or start command for the whole group. For example, you can copy all of the files for a database at the same time.

Perform the following steps to add FlashCopy mappings to a new FlashCopy consistency group:

1. Issue the **svctask mkfcconsistgrp** CLI command to create a FlashCopy consistency group.

The following is an example of the CLI command you can issue to create a FlashCopy consistency group:

```
svctask mkfcconsistgrp -name FCcgrp0 -autodelete
```

Where **FCcgrp0** is the name of the FlashCopy consistency group. The **-autodelete** parameter specifies to delete the consistency group when the last FlashCopy mapping is deleted or removed from the consistency group.

2. Issue the **svcinfolsfconsistgrp** CLI command to display the attributes of the group that you have created.

The following is an example of the CLI command you can issue to display the attributes of a FlashCopy consistency group:

```
svcinfolsfconsistgrp -delim : FCcgrp0
```

The following is an example of the output that is displayed:

```
id:1
name:FCcgrp0
status:idle_or_copied
autodelete:on
FC_mapping_id:0
FC_mapping_name:fcmap0
FC_mapping_id:1
FC_mapping_name:fcmap1
```

Note: For any group that has just been created, the status reported is empty

3. Issue the **svctask chfcmap** CLI command to add FlashCopy mappings to the FlashCopy consistency group:

The following are examples of the CLI commands you can issue to add Flash Copy mappings to the FlashCopy consistency group:

```
svctask chfcmap -consistgrp FCcgrp0 main1copy
svctask chfcmap -consistgrp FCcgrp0 main2copy
```

Where **FCcgrp0** is the name of the FlashCopy consistency group and *main1copy*, *main2copy* are the names of the FlashCopy mappings.

4. Issue the **svcinfolsfmap** CLI command to display the new attributes of the FlashCopy mappings.

The following is an example of the output that is displayed:

```
svcinfolsfmap -delim :
id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:
target_vdisk_name:group_id:group_name:status:progress:copy_rate
0:main1copy:28:maindisk1:29:bkpdisk1:1:FCcgrp0:idle_copied:75
1:main2copy:30:maindisk2:31:bkpdisk2:1:FCcgrp0:idle_copied:50
```

5. Issue the **svcinfolsfconsistgrp** CLI command to display the detailed attributes of the group.

The following is an example of a CLI command that you can issue to display detailed attributes:

```
svcinfolsfconsistgrp -delim : FCcgrp0
```

Where **FCcgrp0** is the name of the FlashCopy consistency group, and **-delim** specifies the delimiter.

The following is an example of the output that is displayed:

```
id:1
name:FCcgrp0
status:idle_or_copied
autodelete:off
FC_mapping_id:0
FC_mapping_name:main1copy
FC_mapping_id:1
FC_mapping_name:main2copy
```

Preparing and starting a FlashCopy consistency group using the CLI

You can use the command-line interface (CLI) to prepare and start a FlashCopy consistency group to start the FlashCopy process.

Successful completion of the FlashCopy process creates a point-in-time copy of the data on the source virtual disk (VDisk) and writes it to the target VDisk (volume) for each mapping in the group. When you have assigned several mappings to a FlashCopy consistency group, you only have to issue a single prepare command to prepare every FlashCopy mapping in the group, and you only have to issue a single start command to start every FlashCopy mapping in the group.

Perform the following steps to prepare and start a FlashCopy consistency group:

1. Issue the **svctask prestartfcconsistgrp** CLI command to prepare the FlashCopy consistency group before the copy process can be started.

Remember: You only have to issue a single prepare command for the whole group to prepare all of the mappings simultaneously.

The following is an example of the CLI command you can issue to prepare the FlashCopy consistency group:

```
svctask prestartfcconsistgrp -restore maintobkpfcopy
```

Where *maintobkpfcopy* is the name of the FlashCopy consistency group.

The optional **restore** parameter forces the consistency group to be prepared even if the target VDisk of one of the mappings in the consistency group is being used as a source VDisk of another active mapping. An active mapping is in the copying, suspended, or stopping state. The group enters the preparing state, and then moves to the prepared state when it is ready.

2. Issue the **svcinfolsfconsistgrp** command to check the status of the FlashCopy consistency group.

The following is an example of the CLI command that you can issue to check the status of the FlashCopy consistency group.

```
svcinfolsfconsistgrp -delim :
```

The following is an example of the output that is displayed:

```
id:name:status
1:maintobkpfcopy:prepared
```

3. Issue the **svctask startfcconsistgrp** CLI command to start the FlashCopy consistency group to make the copy.

Remember: You only have to issue a single start command for the whole group to start all the mappings simultaneously.

The following is an example of the CLI command that you can issue to start the FlashCopy consistency group mappings:

```
svctask startfcconsistgrp -prep -restore maintobkpfcopy
```

Where *maintobkpfcopy* is the name of the FlashCopy consistency group.

When you use the **prep** parameter, the system automatically issues the **prestartfcconsistgrp** command for the group that you specify. When the **restore** parameter is combined with the prep option, you force the consistency group to be started even if the target VDisk of one of the mappings in the consistency group is being used as a source VDisk in another active mapping. An active mapping is in the copying, suspended, or stopping state. The FlashCopy consistency group enters the copying state and returns to the idle_copied state when complete.

4. Issue the **svcinfolsfconsistgrp** command to check the status of the FlashCopy consistency group.

The following is an example of the CLI command that you can issue to check the status of the FlashCopy consistency group:

```
svcinfolsfconsistgrp -delim : maintobkpfcopy
```

Where *maintobkpfcopy* is the name of the FlashCopy consistency group.

The following is an example of the output that is displayed when the process is still copying:

```
id:name:status
1:maintobkpfcopy:copying
```

The following is an example of the output that is displayed when the process has finished copying:

```
id:1
name:maintobkpfcopy
status:idle_copied
autodelete:off
FC_mapping_id:0
FC_mapping_name:main1copy
FC_mapping_id:1
FC_mapping_name:main2copy
```

Stopping a FlashCopy consistency group using the CLI

You can use the command-line interface (CLI) to stop a FlashCopy consistency group.

The **svctask stopfcconsistgrp** CLI command stops all processing that is associated with a FlashCopy consistency group that is in one of the following processing states: prepared, copying, stopping, or suspended.

1. To stop a FlashCopy consistency group, issue the **stopfcconsistgrp** CLI command:

```
svctask stopfcconsistgrp fc_map_id or fc_map_name
```

where *fc_map_id* or *fc_map_name* is the ID or name of the mapping to delete.

2. To stop a consistency group and break the dependency on the source VDisks of any mappings that are also dependent on the target VDisk, issue the following command:

```
svctask stopfcconsistgrp -split fc_map_id or fc_map_name
```

You can specify the **split** parameter when all the maps in the group have a progress of 100. It removes the dependency of any other maps on the source VDisks. You can use this option before you start another FlashCopy consistency group whose target disks are the source disks of the mappings that are being stopped. After the consistency group is stopped with the split option, you can start the other consistency group without the restore option

The command does not return any output.

Deleting a FlashCopy consistency group using the CLI

You can use the command-line interface (CLI) to delete a FlashCopy consistency group.

The **svctask rmfcconsistgrp** CLI command deletes an existing FlashCopy consistency group. The **force** parameter is required only when the consistency group that you want to delete contains mappings.

1. To delete an existing consistency group that does not contain mappings, issue the **svctask rmfcconsistgrp** CLI command:

```
svctask rmfcconsistgrp fc_map_id or fc_map_name
```

where *fc_map_id* or *fc_map_name* is the ID or name of the consistency group to delete.

2. To delete an existing consistency group that contains mappings that are members of the consistency group, issue the following command:

```
svctask rmfcconsistgrp -force fc_map_id or fc_map_name
```

where *fc_map_id* or *fc_map_name* is the ID or name of the mapping to delete.

All the mappings that are associated with the consistency group are removed from the group and changed to stand-alone mappings. To delete a single mapping in the consistency group, you must use the **svctask rmfcmap** command.

The command does not return any output.

Creating Metro Mirror or Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to create Metro Mirror or Global Mirror relationships.

To create Metro Mirror or Global Mirror relationships, perform these steps:

1. To create a Metro Mirror relationship, run the **svctask mkrcrelationship** command. For example, enter:

```
svctask mkrcrelationship -master master_vdisk_id  
-aux aux_vdisk_id -cluster cluster_id
```

Where *master_vdisk_id* is the ID of the master VDisk (volume), *aux_vdisk_id* is the ID of the auxiliary VDisk, and *cluster_id* is the ID of the remote cluster.

2. To create a new Global Mirror relationship, run the **svctask mkrcrelationship** command with the **-global** parameter. For example, enter:

```
svctask mkrcrelationship -master master_vdisk_id  
-aux aux_vdisk_id -cluster cluster_id -global
```

Where *master_vdisk_id* is the ID of the master VDisk, *aux_vdisk_id* is the ID of the auxiliary VDisk, and *cluster_id* is the ID of the remote cluster.

Modifying Metro Mirror or Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to modify certain attributes of Metro Mirror or Global Mirror relationships. You can change only one attribute at a time for each command submission.

To modify Metro Mirror or Global Mirror relationships, run the **svctask chrcrelationship** command.

1. Run the **svctask chrcrelationship** command to change the name of a Metro Mirror or Global Mirror relationship. For example, to change the relationship name, enter:

```
svctask chrcrelationship -name new_rc_rel_name previous_rc_rel_name
```

Where *new_rc_rel_name* is the new name of the relationship and *previous_rc_rel_name* is the previous name of the relationship.

2. Run the **svctask chrcrelationship** command to remove a relationship from whichever consistency group it is a member of. For example, enter:

```
svctask chrcrelationship -force rc_rel_name/id
```

Where *rc_rel_name/id* is the name or ID of the relationship.

Starting and stopping Metro Mirror or Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to start and stop stand-alone Metro Mirror and Global Mirror relationships. Relationships that are members of consistency groups must be started and stopped using the consistency group CLI commands.

To start and stop Metro Mirror or Global Mirror relationships, perform these steps:

1. To start a Metro Mirror or Global Mirror relationship, run the **svctask startrcrelationship** command. For example, enter:

```
svctask startrcrelationship rc_rel_id
```

Where *rc_rel_id* is the ID of the relationship that you want to start in a stand-alone relationship.

2. To stop a Metro Mirror or Global Mirror relationship, run the **svctask stopprcrelationship** command. This command applies to a stand-alone relationship.

For example, enter:

```
svctask stopprcrelationship rc_rel_id
```

Where *rc_rel_id* is the ID of the stand-alone relationship that you want to stop mirroring I/O.

Displaying the progress of Metro Mirror or Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to display the background copy of Metro Mirror or Global Mirror relationships as a percentage. When the initial background copy process for a relationship has completed, null is displayed for the progress of that relationship.

To display the progress of the background copy of Metro Mirror or Global Mirror relationships, run the **lsrrelationshipprogress** command.

1. To display data progress without headings for columns of data or for each item of data in a Metro Mirror or Global Mirror relationship, run the **lsrrelationshipprogress -nohdr** command. For example, to display data of the relationship with headings suppressed, enter:

```
lsrrelationshipprogress -nohdr rc_rel_name
```

Where *rc_rel_name* is the name of the specified object type.

2. To display the progress of a background copy of a Metro Mirror or Global Mirror relationship as a percentage, run the **lsrrelationshipprogress -delim** command. The colon character (:) separates all items of data in a concise view, and the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter. For example, enter:

```
svcinfolsrrelationshipprogress -delim : 0
```

The resulting output is displayed, such as in this example:

```
id:progress
0:58
```

Switching Metro Mirror or Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to reverse the roles of primary and secondary virtual disks (volumes) in a stand-alone Metro Mirror or Global Mirror relationship when that relationship is in a consistent state. Relationships that are members of consistency groups must be switched by using the consistency group CLI commands.

To switch the roles of primary and secondary VDisks in Metro Mirror or Global Mirror relationships, follow these steps:

1. To make the master disk in a Metro Mirror or Global Mirror relationship to be the primary, run the **svctask switchrcrelationship -primary master** command. For example, enter:

```
svctask switchrcrelationship -primary master rc_rel_id
```

Where *rc_rel_id* is the ID of the relationship to switch.

2. To make the auxiliary disk in a Metro Mirror or Global Mirror relationship to be the primary, run the **svctask switchrcrelationship -primary aux** command. For example, enter:

```
svctask switchrcrelationship -primary aux rc_rel_id
```

Where *rc_rel_id* is the ID of the relationship to switch.

Deleting Metro Mirror and Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to delete Metro Mirror and Global Mirror relationships.

To delete Metro Mirror and Global Mirror relationships, run the **svctask rmrcrelationship** command. For example, enter:

```
svctask rmrcrelationship rc_rel_name/id
```

where *rc_rel_name/id* is the name or ID of the relationship.

Creating Metro Mirror or Global Mirror consistency groups using the CLI

You can use the command-line interface (CLI) to create Metro Mirror or Global Mirror consistency groups.

To create Metro Mirror or Global Mirror consistency groups, perform these steps:

1. To create a Metro Mirror or Global Mirror consistency group, run the **svctask mkrconsistgrp** command. For example, enter:

```
svctask mkrconsistgrp -name new_name -cluster cluster_id
```

where *new_name* is the name of the new consistency group and *cluster_id* is the ID of the remote cluster for the new consistency group. If **-cluster** is not specified, a consistency group is created only on the local cluster. The new consistency group does not contain any relationships and will be in the empty state.

2. To add Metro Mirror or Global Mirror relationships to the group, run the **svctask chrcrelationship** command. For example, enter:

```
svctask chrcrelationship -consistgrp consist_group_name rc_rel_id
```

where *consist_group_name* is the name of the new consistency group to assign the relationship to and *rc_rel_id* is the ID of the relationship.

Modifying Metro Mirror or Global Mirror consistency groups using the CLI

You can use the command-line interface (CLI) to assign a new name or modify the name of an existing Metro Mirror or Global Mirror consistency group.

To assign or modify the name of a Metro Mirror or Global Mirror consistency group, run the **svctask chrconsistgrp** command.

1. Run the **svctask chrconsistgrp** command to assign a new name of a Metro Mirror or Global Mirror consistency group. For example, enter:

```
svctask chrconsistgrp -name new_name_arg
```

Where *new_name_arg* is the assigned new name of the consistency group.

2. Run the **svctask chrconsistgrp** command to change the name of the consistency group. For example, enter:

```
svctask chrconsistgrp -name new_consist_group_name previous_consist_group_name
```

Where *new_consist_group_name* is the assigned new name of the consistency group and *previous_consist_group_name* is the previous name of the consistency group.

Starting and stopping Metro Mirror or Global Mirror consistency-group copy processes using the CLI

You can use the command-line interface (CLI) to start and stop Metro Mirror or Global Mirror consistency-group copy processes.

To start and stop Metro Mirror or Global Mirror consistency-group copy processes, perform these steps:

1. To start a Metro Mirror or Global Mirror consistency-group copy process, set the direction of copy if it is undefined and optionally mark the secondary VDIs of the consistency group as clean. Run the **svctask startrcconsistgrp** command. For example, enter:

```
svctask startrcconsistgrp rc_consist_group_id
```

Where *rc_consist_group_id* is the ID of the consistency group to start processing.

2. To stop the copy process for a Metro Mirror or Global Mirror consistency group, run the **svctask stoprcconsistgrp** command.

For example, enter:

```
svctask stoprconsistgrp rc_consist_group_id
```

Where *rc_consist_group_id* is the ID of the consistency group that you want to stop processing.

If the group is in a consistent state, you can also use this command to enable write access to the secondary virtual disks (VDisks) in the group.

Deleting Metro Mirror or Global Mirror consistency groups using the CLI

You can use the command-line interface (CLI) to delete Metro Mirror or Global Mirror consistency groups.

To delete existing Metro Mirror or Global Mirror consistency groups, follow these steps:

1. To delete a Metro Mirror or Global Mirror consistency group, run the **svctask rmrconsistgrp** command. For example, enter:

```
svctask rmrconsistgrp rc_consist_group_id
```

Where *rc_consist_group_id* is the ID of the consistency group to delete.

2. If a Metro Mirror or Global Mirror consistency group is not empty, you must use the **-force** parameter to delete the consistency group. For example, enter:

```
svctask rmrconsistgrp -force rc_consist_group_id
```

Where *rc_consist_group_id* is the ID of the consistency group to delete. This command causes all relationships that are members of the deleted group to become stand-alone relationships.

Creating Metro Mirror and Global Mirror partnerships using the CLI

You can use the command-line interface (CLI) to create Metro Mirror and Global Mirror partnerships between two clusters.

Perform the following steps to create Metro Mirror and Global Mirror partnerships:

1. To create Metro Mirror and Global Mirror partnerships, run the **svctask mkpartnership** command. For example, enter:

```
svctask mkpartnership -bandwidth bandwidth_in_mbps remote_cluster_id
```

where *bandwidth_in_mbps* specifies the bandwidth (in megabytes per second) that is used by the background copy process between the clusters and *remote_cluster_id* is the ID of the remote cluster.

2. Run the **svctask mkpartnership** command from the remote cluster. For example, enter:

```
svctask mkpartnership -bandwidth bandwidth_in_mbps local_cluster_id
```

where *bandwidth_in_mbps* specifies the bandwidth (in megabytes per second) that is used by the background copy process between the clusters and *local_cluster_id* is the ID of the local cluster.

Modifying Metro Mirror and Global Mirror partnerships using the CLI

You can use the command-line interface (CLI) to modify Metro Mirror and Global Mirror partnerships.

- | The partnership bandwidth, which is also known as *background copy*, controls the rate at which data is sent from the local cluster to the remote cluster. The partnership bandwidth can be changed to help manage the use of intercluster links. It is measured in megabytes per second (MBps).

Perform the following steps to modify Metro Mirror and Global Mirror partnerships:

1. To modify Metro Mirror and Global Mirror partnerships, run the **svctask chpartnership** command. For example, enter:

```
svctask chpartnership -bandwidth bandwidth_in_mbps remote_cluster_id
```

where *bandwidth_in_mbps* is the new bandwidth (in megabytes per second) from the local cluster to the remote cluster, and *remote_cluster_id* is the ID of the remote cluster.

2. Run the **svctask chpartnership** command from the remote cluster. For example, enter:

```
svctask chpartnership -bandwidth bandwidth_in_mbps local_cluster_id
```

where *bandwidth_in_mbps* is the new bandwidth (in megabytes per second) from the remote cluster to the local cluster, and *local_cluster_id* is the ID of the local cluster.

Starting and stopping Metro Mirror and Global Mirror partnerships using the CLI

You can use the command-line interface (CLI) to start and stop Metro Mirror and Global Mirror partnerships.

Perform the following steps to start and stop Metro Mirror and Global Mirror partnerships:

1. To start a Metro Mirror or Global Mirror partnership, run the **svctask chpartnership** command from either cluster. For example, enter:

```
svctask chpartnership -start remote_cluster_id
```

Where *remote_cluster_id* is the ID of the remote cluster. The **svctask mkpartnership** command starts the partnership by default.
2. To stop a Metro Mirror or Global Mirror partnership, run the **svctask chpartnership** command from either cluster.
For example, enter:

```
svctask chpartnership -stop remote_cluster_id
```

Where *remote_cluster_id* is the ID of the remote cluster.

Deleting Metro Mirror and Global Mirror partnerships using the CLI

You can use the command-line interface (CLI) to delete Metro Mirror and Global Mirror partnerships.

Perform the following steps to delete Metro Mirror and Global Mirror partnerships:

1. If a Metro Mirror or Global Mirror partnership has configured relationships or groups, you must stop the partnership before you can delete it. For example, enter:

```
svctask chpartnership -stop remote_cluster_id
```

Where *remote_cluster_id* is the ID of the remote cluster.
2. To delete a Metro Mirror and Global Mirror partnership, run the **svctask rmpartnership** command from either cluster. For example, enter:

```
svctask rmpartnership remote_cluster_id
```

Where *remote_cluster_id* is the ID of the remote cluster.

Determining the WWPNs of a node using the CLI

You can determine the worldwide port names (WWPNs) of a node using the command-line interface (CLI).

Perform the following steps to determine the WWPNs of a node:

1. Issue the **svcinfolnode** CLI command to list the nodes in the cluster.
2. Record the name or ID of the node for which you want to determine the WWPNs.
3. Issue the **svcinfolnode** CLI command and specify the node name or ID that was recorded in step 2.
The following is an example of the CLI command you can issue:

```
svcinfolnode node1
```

Where *node1* is the name of the node for which you want to determine the WWPNs.
4. Record the four port IDs (WWPNs).

Listing node-dependent VDisks (volumes) using the CLI

You can use the command-line interface (CLI) to list the virtual disks (VDisks) that are dependent on the status of a node.

If a node goes offline or is removed from a cluster, all VDisks (volumes) that are dependent on the node go offline. Before taking a node offline or removing a node from a cluster, run the **lsdependentvdisks** command to identify any node-dependent VDisks.

By default, the **lsdependentvdisks** command also checks all available quorum disks. If the quorum disks are accessible only through the specified node, the command returns an error.

Various scenarios can produce node-dependent volumes. The following examples are common scenarios in which the **lsnode-dependentvdisks** command will return node-dependent volumes:

1. The node contains solid-state drives (SSDs) and also contains the only synchronized copy of a mirrored volume.
2. The node is the only node that can access an MDisk on the SAN fabric.
3. The other node in the I/O group is offline (all volumes in the I/O group are returned).
4. Pinned data in the cache is stopping the partner node from joining the I/O group.

To resolve (1), allow volume mirror synchronizations between SSD MDisks to complete. To resolve (2-4), bring any offline MDisks online and repair any degraded paths.

Note: The command lists the node-dependent VDisks at the time the command is run; subsequent changes to a cluster require running the command again.

1. Issue the **svcinfoldependentvdisks** CLI command.

The following example shows the CLI format for listing the VDisks that are dependent on node01:

```
svcinfoldependentvdisks -drive -delim : 0:1
```

The following example shows the output that is displayed:

```
vdisk_id:vdisk_name  
4:vdisk4  
5:vdisk5
```

2. If the **svcinfoldependentvdisks** command returns an error, you must move your quorum disks to MDisks that are accessible through all nodes. Rerun the command until no errors are returned.
3. Reissue the **svcinfoldependentvdisks** command. When the command returns no VDisks, the cluster is free from any node-dependent VDisks.

The following example shows the command syntax for listing the VDisks that are dependent on node01:

```
svcinfoldependentvdisks -node01 :
```

The following example shows the command output if there are no node-dependent VDisks in the cluster:

```
vdisk_id      vdisk_name
```

Determining the VDisk name from the device identifier on the host

You can use the command-line interface (CLI) to determine the virtual disk (VDisk) name from the device identifier on the host.

Each VDisk that is exported by the SAN Volume Controller is assigned a unique device identifier. The device identifier uniquely identifies the VDisk (volume) and can be used to determine which VDisk corresponds to the volume that the host sees.

Perform the following steps to determine the VDisk name from the device identifier:

1. Find the device identifier. For example, if you are using the subsystem device driver (SDD), the disk identifier is referred to as the virtual path (vpath) number. You can issue the following SDD command to find the vpath serial number:

```
datapath query device
```

For other multipathing drivers, refer to the documentation that is provided with your multipathing driver to determine the device identifier.

2. Find the host object that is defined to the SAN Volume Controller and corresponds with the host that you are working with.

- a. Find the worldwide port numbers (WWPNs) by looking at the device definitions that are stored by your operating system. For example, on AIX the WWPNs are in the ODM and if you use Windows you have to go into the HBA Bios.

- b. Verify which host object is defined to the SAN Volume Controller for which these ports belong. The ports are stored as part of the detailed view, so you must list each host by issuing the following CLI command:

```
svcinfo lshost id | name
```

Where *name/id* is the name or ID of the host.

- c. Check for matching WWPNs.

3. Issue the following command to list the VDisk-to-host mappings:

```
svcinfo lshostvdiskmap hostname
```

Where *hostname* is the name of the host.

4. Find the VDisk UID that matches the device identifier and record the VDisk name or ID.

Determining the host that a VDisk (volume) is mapped to

You can determine the host that a virtual disk (VDisk) is mapped to using the command-line interface (CLI). To view the host mapping for a volume in the management GUI, select **Volumes** → **Volumes by Hosts**.

Perform the following steps to determine the host that the VDisk (volume) is mapped to:

1. Find the VDisk name or ID that you want to check.
2. Issue the following CLI command to list the hosts that this VDisk is mapped:

```
svcinfo lsvdiskhostmap vdiskname/id
```

where *vdiskname/id* is the name or ID of the VDisk.

3. Find the host name or ID to determine which host this VDisk is mapped to.
 - If no data is returned, the VDisk is not mapped to any hosts.

Determining the relationship between VDIsks (volumes) and MDisks using the CLI

You can determine the relationship between virtual disks (VDisks) and managed disks (MDisks) using the command-line interface (CLI).

Select one or more of the following options to determine the relationship between VDIsks (volumes) and MDisks:

- To display a list of the IDs that correspond to the MDisks that comprise the VDisk, issue the following CLI command:

```
svcinfo lsvdiskmember vdiskname/id
```

where *vdiskname/id* is the name or ID of the VDisk.

- To display a list of IDs that correspond to the VDIsks that are using this MDisk, issue the following CLI command:

```
svcinfolsmdiskmember mdiskname/id
```

where *mdiskname/id* is the name or ID of the MDisk.

- To display a table of VDisk IDs and the corresponding number of extents that are being used by each VDisk, issue the following CLI command:

```
svcinfolsmdiskextent mdiskname/id
```

where *mdiskname/id* is the name or ID of the MDisk.

- To display a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the given VDisk, issue the following CLI command:

```
svcinfolsvdiskextent vdiskname/id
```

where *vdiskname/id* is the name or ID of the VDisk.

Determining the relationship between MDisks and controller LUNs using the CLI

You can determine the relationship between managed disks (MDisks) and RAID arrays or LUNs using the command-line interface (CLI).

Each MDisk corresponds with a single RAID array, or with a single partition on a given RAID array. Each RAID controller defines a LUN number for this disk. The LUN number and controller name or ID are needed to determine the relationship between MDisks and RAID arrays or partitions.

Perform the following steps to determine the relationship between MDisks and RAID arrays:

1. Issue the following command to display a detailed view of the MDisk:

```
svcinfolsmdisk mdiskname
```

Where *mdiskname* is the name of the MDisk for which you want to display a detailed view.

2. Record the controller name or controller ID and the controller LUN number.
3. Issue the following command to display a detailed view of the controller:

```
svcinfolcontroller controllername
```

Where *controllername* is the name of the controller that you recorded in step 2.

4. Record the vendor ID, product ID, and WWNN. You can use this information to determine what is being presented to the MDisk.
5. From the native user interface for the given controller, list the LUNs it is presenting and match the LUN number with that noted in step 1. This tells you the exact RAID array or partition that corresponds with the MDisk.

Increasing the size of your cluster using the CLI

You can increase throughput by adding more nodes to the cluster. The nodes must be added in pairs and assigned to a new I/O group.

Perform the following steps to increase the size of your cluster:

1. Add a node to your cluster and repeat this step for the second node.
2. If you want to balance the load between the existing I/O groups and the new I/O groups, you can migrate your virtual disks (VDisks) to new I/O groups. Repeat this step for all VDIs (volumes) that you want to assign to the new I/O group.

Adding a node to increase the size of a cluster using the CLI

You can use the command-line interface (CLI) to increase the size of a cluster by adding two nodes to create a new I/O group.

Attention: If you are adding a node that was previously removed from a cluster, ensure that either the following two conditions have been met:

- The WWPN for the removed node is swapped with the node that replaces it.
- All hosts that accessed the removed node through its WWPNs have been reconfigured to use the WWPN for the new node.

Failure to do either of these action can result in data corruption.

Complete the following steps to add a node and increase the size of a cluster:

1. Install the new nodes and connect to the fibre channel.
2. Issue the following command to verify that the node is detected on the fabric:
`svcinfolsnodecandidate`
3. Using the front panel of the node, record the WWNN.
4. Issue the following command to determine the I/O group where the node should be added:
`svcinfolsiogrp`
5. Record the name or ID of the first I/O group that has a node count of zero (0). You will need the ID for the next step.

Note: You only need to do this step for the first node that is added. The second node of the pair uses the same I/O group number.

6. Record the following information for future reference:
 - Node serial number.
 - Worldwide node name.
 - All of the worldwide port names.
 - The name or ID of the I/O group that contains the node.
7. Issue the following command to add the node to the cluster:

```
svctask addnode -wwnodename WWNN -iogrp newiogrpname/id [-name newnodename]
```

Where *WWNN* is the WWNN of the node, *newiogrpname/id* is the name or ID of the I/O group that you want to add the node to and *newnodename* is the name that you want to assign to the node. If you do not specify a new node name, a default name is assigned; however, it is recommended you specify a meaningful name.

8. Record the following information for future reference:
 - Node serial number.
 - Worldwide node name.
 - All of the worldwide port names.
 - The name or ID of the I/O group that contains the node.
9. Issue the following command to verify that the node is online:

```
svcinfolsnode
```

You may need to reconfigure your storage systems to allow the new I/O group nodes to access them. If the disk controller uses mapping to present RAID arrays or partitions to the cluster and the WWNNs or the worldwide port names have changed, you must modify the port groups that belong to the cluster.

Migrating a VDisk to a new I/O group using the CLI

You can use the command-line interface (CLI) to migrate a virtual disk (VDisk) to a new I/O group to increase the size of your cluster.

- | You can migrate a VDisk (volume) to a new I/O group to manually balance the workload across the nodes in the cluster. However, you might end up with a pair of nodes that are overworked and another pair that are not worked. Follow this procedure to migrate a single VDisk to a new I/O group. Repeat for other VDIs as required.

Attention: This is a disruptive procedure. Access to the VDisk is lost while you follow this procedure. Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

Perform the following steps to migrate a single VDisk:

1. Quiesce all I/O operations for the VDisk. You might have to determine the hosts that are using this VDisk in advance.
2. Before migrating the VDisk, it is essential that for each device identifier that is presented by the VDisk you intend to move, the subsystem device driver (SDD) or other multipathing driver configuration is updated to remove the device identifiers. Failure to do this can result in data corruption. See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* or the documentation that is provided with your multipathing driver for details about how to dynamically reconfigure device identifiers for the given host operating system.
3. Issue the following command to check if the VDisk is part of a relationship or mapping:

```
svcinfolsvdisk vdiskname/id
```

Where *vdiskname/id* is the name or ID of the VDisk.

- a. Find the **FC_id** and **RC_id** fields. If these are not blank, the VDisk is part of a mapping or relationship.
 - b. Stop or delete any FlashCopy mappings, Global Mirror, or Metro Mirror relationships that use this VDisk.
4. Issue the following command to migrate the VDisk:

```
svctask chvdisk -iogrp newiogrpname/id -node preferred_node vdiskname/id
```

where *preferred_node* is the name of the node that you want to move the VDisk, *newiogrpname/id* is the name or ID of the I/O group where you want to migrate the VDisk and *vdiskname/id* is the name or ID of the VDisk that you want to migrate.
 5. Discover the new device identifiers and check that each device identifier presents the correct number of paths. See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* or the documentation that is provided with your multipathing driver for details about how to discover device identifiers for the given host operating system.

Validating and repairing mirrored VDisk copies using the CLI

- | You can use the **repairvdiskcopy** command from the command-line interface (CLI) to validate and repair mirrored VDisk (volume) copies.

- | **Attention:** Run the **repairvdiskcopy** command only if all VDisk (volume) copies are synchronized.

When you issue the **repairvdiskcopy** command, you must use only one of the **-validate**, **-medium**, or **-resync** parameters. You must also specify the name or ID of the VDisk to be validated and repaired as the last entry on the command line. After you issue the command, no output is displayed.

-validate

Use this parameter if you only want to verify that the mirrored VDisk copies are identical. If any difference is found, the command stops and logs an error that includes the logical block address (LBA) and the length of the first difference. You can use this parameter, starting at a different LBA each time to count the number of differences on a VDisk.

-medium

Use this parameter to convert sectors on all VDisk copies that contain different contents into virtual

medium errors. Upon completion, the command logs an event, which indicates the number of differences that were found, the number that were converted into medium errors, and the number that were not converted. Use this option if you are unsure what the correct data is, and you do not want an incorrect version of the data to be used.

-resync

Use this parameter to overwrite contents from the specified primary VDisk copy to the other VDisk copy. The command corrects any differing sectors by copying the sectors from the primary copy to the copies being compared. Upon completion, the command process logs an event, which indicates the number of differences that were corrected. Use this action if you are sure that either the primary VDisk copy data is correct or that your host applications can handle incorrect data.

-startlba *lba*

Optionally, use this parameter to specify the starting Logical Block Address (LBA) from which to start the validation and repair. If you previously used the **validate** parameter, an error was logged with the LBA where the first difference, if any, was found. Reissue **repairvdiskcopy** with that LBA to avoid reprocessing the initial sectors that compared identically. Continue to reissue **repairvdiskcopy** using this parameter to list all the differences.

Issue the following command to validate and, if necessary, automatically repair mirrored copies of the specified VDisk:

```
svctask repairvdiskcopy -resync -startlba 20 vdisk8
```

Notes:

1. Only one **repairvdiskcopy** command can run on a VDisk at a time.
2. Once you start the **repairvdiskcopy** command, you cannot use the command to stop processing.
3. The primary copy of a mirrored VDisk cannot be changed while the **repairvdiskcopy -resync** command is running.
4. If there is only one mirrored copy, the command returns immediately with an error.
5. If a copy being compared goes offline, the command is halted with an error. The command is not automatically resumed when the copy is brought back online.
6. In the case where one copy is readable but the other copy has a medium error, the command process automatically attempts to fix the medium error by writing the read data from the other copy.
7. If no differing sectors are found during **repairvdiskcopy** processing, an informational error is logged at the end of the process.

Checking the progress of validation and repair of VDisk copies using the CLI

Use the **lsrepairvdiskcopyprogress** command to display the progress of mirrored VDisk validation and repairs. You can specify a VDisk copy using the **-copy *id*** parameter. To display the VDIsks that have two or more copies with an active task, specify the command with no parameters; it is not possible to have only one VDisk copy with an active task.

To check the progress of validation and repair of mirrored VDIsks, issue the following command:

```
svcinfo lsrepairvdiskcopyprogress -delim :
```

The following example shows how the command output is displayed:

```
vdisk_id:vdisk_name:copy_id:task:progress:estimated_completion_time
0:vdisk0:0:medium:50:070301120000
0:vdisk0:1:medium:50:070301120000
```

Repairing a space-efficient VDisk using the CLI

You can use the **repairsevdiskcopy** command from the command-line interface to repair the metadata on a space-efficient virtual disk (VDisk).

The **repairsevdiskcopy** command automatically detects and repairs corrupted metadata. The command holds the VDisk (volume) offline during the repair, but does not prevent the disk from being moved between I/O groups.

If a repair operation completes successfully and the volume was previously offline because of corrupted metadata, the command brings the volume back online. The only limit on the number of concurrent repair operations is the number of virtual disk copies in the configuration.

When you issue the **repairsevdiskcopy** command, you must specify the name or ID of the VDisk to be repaired as the last entry on the command line. Once started, a repair operation cannot be paused or cancelled; the repair can only be terminated by deleting the copy.

Attention: Use this command only to repair a space-efficient VDisk (thin-provisioned volume) that has reported corrupt metadata.

Issue the following command to repair the metadata on a space-efficient VDisk:

```
svctask repairsevdiskcopy vdisk8
```

After you issue the command, no output is displayed.

Notes:

1. Because the volume is offline to the host, any I/O that is submitted to the volume while it is being repaired fails.
2. When the repair operation completes successfully, the corrupted metadata error is marked as fixed.
3. If the repair operation fails, the volume is held offline and an error is logged.

Checking the progress of the repair of a space-efficient VDisk using the CLI

Issue the **lsrepairsevdiskcopyprogress** command to list the repair progress for space-efficient VDisk copies of the specified VDisk. If you do not specify a VDisk, the command lists the repair progress for all space-efficient copies in the cluster.

Note: Only run this command after you run the **svctask repairsevdiskcopy** command, which you must only run as required by the Directed Maintenance Procedures or by IBM support.

Recovering from offline VDIsks using the CLI

If a node or an I/O group fails, you can use the command-line interface (CLI) to recover offline virtual disks (volumes).

If you have lost both nodes in an I/O group and have, therefore, lost access to all the VDIsks (volumes) that are associated with the I/O group, you must perform one of the following procedures to regain access to your volumes. Depending on the failure type, you might have lost data that was cached for these volumes and the volumes are now offline.

Data loss scenario 1

One node in an I/O group has failed and failover has started on the second node. During the failover process, the second node in the I/O group fails before the data in the write cache is written to hard disk. The first node is successfully repaired but its hardened data is not the most recent version that is

committed to the data store; therefore, it cannot be used. The second node is repaired or replaced and has lost its hardened data, therefore, the node has no way of recognizing that it is part of the cluster.

Perform the following steps to recover from an offline VDisk when one node has down-level hardened data and the other node has lost hardened data:

1. Recover the node and add it back into the cluster.
2. Delete all IBM FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the offline volumes.
3. Run the **recovervdisk**, **recovervdiskbyiogrp** or **recovervdiskbycluster** command.
4. Re-create all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the volumes.

Data loss scenario 2

Both nodes in the I/O group have failed and have been repaired. The nodes have lost their hardened data, therefore, the nodes have no way of recognizing that they are part of the cluster.

Perform the following steps to recover from an offline volume when both nodes have lost their hardened data and cannot be recognized by the cluster:

1. Delete all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the offline volumes.
2. Run the **recovervdisk**, **recovervdiskbyiogrp** or **recovervdiskbycluster** command.
3. Recreate all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the volumes.

Recovering a node and returning it to the cluster using the CLI

After a node or an I/O group fails, you can use the command-line interface (CLI) to recover a node and return it to the cluster.

Perform the following steps to recover a node and return it to the cluster:

1. Issue the following command to verify that the node is offline:
`svcinfolnode`
2. Issue the following command to remove the old instance of the offline node from the cluster:
`svctask rmnode nodename/id`

Where *nodename/id* is the name or ID of the node.

3. Issue the following command to verify that the node can be seen on the fabric:
`svcinfolnodecandidate`

Note: Remember the worldwide node names (WWNNs) for each node because you will need them in the following step.

4. If the nodes are repaired by replacing the service controller, or the node is replaced, be sure to follow the replacement instructions for the specific node or controller. You will be instructed to reset the WWNN of the node to that of the original node. If you do not do that, you may need to reconfigure your SAN fabric, your hosts, and your storage systems.

Attention: If more than one I/O group is affected, ensure that you are adding the node to the same I/O group from which it was removed. Failure to do this can result in data corruption. Use the information that was recorded when the node was originally added to the cluster. This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster. If you do not have access to this information, call the IBM Support Center to add the node back into the cluster without corrupting the data. If you are adding the node into the cluster for the first time, you must record the following information:

- Node serial number
- WWNN
- All WWPNNs
- I/O group that contains the node

5. Issue the following command to add the node back into the cluster:

```
svctask addnode -wwnname WWNN -iogrp
IOGRPNAME/ID [-name NODENAME]
```

Where *WWNN* is the worldwide node name, *IOGRPNAME/ID* is the I/O group name or ID and *NODENAME* is the name of the node.

In a service situation, a node should normally be added back into a cluster using the original node name. As long as the partner node in the I/O group has not been deleted too, this is the default name used if **-name** is not specified.

6. Issue the following command to verify that the node is online:

```
svcinfolnode
```

Recovering offline VDisks using the CLI

You can recover offline virtual disks (VDisks) using the command-line interface (CLI).

1. Perform the following steps to recover offline VDisks (volumes):

1. Issue the following CLI command to list all VDisks that are offline and belong to an I/O group, enter:

```
svcinfolsvdisk -filtervalue IO_group_name=
IOGRPNAME/ID:status=offline
```

where *IOGRPNAME/ID* is the name of the I/O group that failed.

2. To acknowledge data loss for a VDisk with a *fast_write_state* of **corrupt** and bring the VDisk back online, enter:

```
svctask recovervdisk vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the VDisk.

Notes:

- If the specified VDisk is space-efficient or has space-efficient copies, the **recovervdisk** command starts the space-efficient repair process.
- If the specified VDisk is mirrored, the **recovervdisk** command starts the resynchronization process.

3. To acknowledge data loss for all virtual disks in an I/O group with a *fast_write_state* of **corrupt** and bring them back online, enter:

```
svctask recovervdiskbyiogrp io_group_id | io_group_name
```

where *io_group_id* | *io_group_name* is the name or ID of the I/O group.

Notes:

- If any VDisk is space-efficient or has space-efficient copies, the **recovervdiskbyiogrp** command starts the space-efficient repair process.
- If any VDisk is mirrored, the **recovervdiskbyiogrp** command starts the resynchronization process.

4. To acknowledge data loss for all VDIs in the cluster with a *fast_write_state* of **corrupt** and bring them back online, enter:

```
svctask recovervdiskbycluster
```

Notes:

- If any VDisk is space-efficient or has space-efficient copies, the **recovervdiskbycluster** command starts the space-efficient repair process.
- If any VDisk is mirrored, the **recovervdiskbycluster** command starts the resynchronization process.

Moving offline VDIs to their original I/O group using the CLI

You can move offline virtual disks (VDIs) to their original I/O group using the command-line interface (CLI).

After a node or an I/O group fails, you can use the following procedure to move offline VDIs (volumes) to their original I/O group.

Attention: Do not move VDIs to an offline I/O group. Ensure that the I/O group is online before you move the VDIs back to avoid any further data loss.

Perform the following steps to move offline VDIs to their original I/O group:

1. Issue the following command to move the VDisk back into the original I/O group:

```
svctask chvdisk -iogrp IOGRPNAME/ID -force  
vdiskname/ID
```

where *IOGRPNAME/ID* is the name or ID of the original I/O group and *vdiskname/ID* is the name or ID of the offline VDisk.

2. Issue the following command to verify that the VDIs are now online:

```
svcinfolsvdisk -filtervalue IO_group_name=  
IOGRPNAME/ID
```

where *IOGRPNAME/ID* is the name or ID of the original I/O group.

Recording WWPN changes of replaced host HBAs

You can use the command-line interface (CLI) to record a change to a defined host object.

Because it is sometimes necessary to replace the host-bus adapter (HBA) that connects the host to the SAN, you must inform the SAN Volume Controller of the new worldwide port names (WWPNs) that this HBA contains.

Ensure that your switch is zoned correctly.

Perform the following steps to inform the SAN Volume Controller of a change to a defined host object:

1. Issue the following CLI command to list the candidate HBA ports:

```
svcinfolshbaportcandidate
```

You should see a list of the HBA ports that are available for addition to host objects. One or more of these HBA ports should correspond with the one or more WWPNs that belong to the new HBA port.

2. Locate the host object that corresponds with the host in which you have replaced the HBA. The following CLI command lists all the defined host objects:

```
svcinfolshost
```

3. Issue the following CLI command to list the WWPNs that are currently assigned to the host object:

```
svcinfolshost hostobjectname
```

where *hostobjectname* is the name of the host object.

4. Issue the following CLI command to add the new ports to the existing host object:

```
svctask addhostport -hbawwpn one or more existing WWPNS  
separated by : hostobjectname/ID
```

where *one or more existing WWPNS separated by :* is the WWPNS that are currently assigned to the host object and *hostobjectname/ID* is the name or ID of the host object.

5. Issue the following CLI command to remove the old ports from the host object:

```
svctask rmhostport -hbawwpn one or more existing WWPNS  
separated by : hostobjectname/ID
```

where *one or more existing WWPNS separated by :* is the WWPNS that are currently assigned to the host object and *hostobjectname/ID* is the name or ID of the host object.

Any mappings that exist between the host object and the virtual disks (VDisks) are automatically applied to the new WWPNS. Therefore, the host sees the VDisks as the same SCSI LUNs as before.

See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* or the documentation that is provided with your multipathing driver for additional information about dynamic reconfiguration.

Expanding VDisks (volumes) using the CLI

- | You can use the command-line interface (CLI) to expand a virtual disk (VDisk).
- | VDisks (volumes) that are mapped for FlashCopy or that are in Metro Mirror relationships cannot be expanded.

Ensure that you have run Windows Update and have applied all recommended updates to your system before you attempt to expand a VDisk that is mapped to a Windows host.

Determine the exact size of the source or master VDisk by issuing the following CLI command:

```
svcinfolsvdisk -bytes vdiskname
```

where *vdiskname* is the name of the VDisk for which you want to determine the exact size.

VDisks can be expanded under Windows concurrently with I/O operations.

You can expand VDisks for the following reasons:

- To increase the available capacity on a particular VDisk that is already mapped to a host.
- To increase the size of a VDisk so that it matches the size of the source or master VDisk and so that it can be used in a FlashCopy mapping or Metro Mirror relationship.

A VDisk that is not mapped to any hosts and does not contain any data can be expanded at any time. If the VDisk contains data that is in use, you can expand the VDisks if your host has a supported AIX or Microsoft Windows operating system.

The following table provides the supported operating systems and requirements for expanding VDisks that contain data:

Operating system	Supported	Requirement
AIX	Yes	AIX version 5.2 or later
HP-UX	No	-
Linux	No	-
SUN Solaris	No	-
Microsoft Windows NT	No	-

Operating system	Supported	Requirement
Microsoft Windows	Yes	Windows version 2000 or later

Expanding a VDisk (volume) that is mapped to an AIX host

The SAN Volume Controller supports the ability to dynamically expand the size of a virtual disk (VDisk) if the AIX host is using AIX version 5.2 or later.

The **chvg** command options provide the ability to expand the size of a physical volume that the Logical Volume Manager (LVM) uses, without interruptions to the use or availability of the system. See the *AIX System Management Guide Operating System and Devices* for more information.

Expanding a VDisk that is mapped to a Microsoft Windows host using the CLI

You can use the command-line interface (CLI) to dynamically expand the size of a virtual disk (VDisk) that is mapped to a Microsoft Windows host.

1 Perform the following steps to expand a VDisk (volume) that is mapped to a Windows host:

1. Issue the following CLI command to expand the VDisk:

```
svctask expandvdisksize -size disk_size -unit
b | kb | mb | gb | tb | pb vdisk_name/vdisk_id
```

where *disk_size* is the capacity by which you want to expand the VDisk, *b | kb | mb | gb | tb | pb* is the data unit to use in conjunction with the capacity and *vdisk_name/vdisk_id* is the name of the VDisk or the ID of the VDisk to expand.

2. On the Windows host, start the Computer Management application and open the Disk Management window under the Storage branch.

You will see the VDisk that you expanded now has some unallocated space at the end of the disk.

You can expand dynamic disks without stopping I/O operations in most cases. However, in some applications the operating system might report I/O errors. When this problem occurs, either of the following entries might be recorded in the System event log:

```
Event Type: Information
Event Source: dmio
Event Category: None
Event ID: 31
Description: dmio:
Harddisk0 write error at block ##### due to
disk removal
```

```
Event Type: Information
Event Source: dmio
Event Category: None
Event ID: 34
Description: dmio:
Harddisk0 is re-online by PnP
```

Attention: This is a known problem with Windows 2000 and is documented in the Microsoft knowledge base as article Q327020. If either of these errors are seen, run Windows Update and apply the recommended fixes to resolve the problem.

If the Computer Management application was open before you expanded the VDisk, use the Computer Management application to issue a rescan command.

If the disk is a Windows basic disk, you can create a new primary or extended partition from the unallocated space.

If the disk is a Windows dynamic disk, you can use the unallocated space to create a new volume (simple, striped, mirrored) or add it to an existing volume.

Shrinking a virtual disk using the CLI

- | You can reduce the size of a virtual disk (VDisk) (volume) using the command-line interface (CLI).
- | VDIsks (volumes) can be reduced in size, if it is necessary. You can make a target or auxiliary VDisk the same size as the source or master VDisk when you create FlashCopy® mappings, Metro Mirror relationships, or Global Mirror relationships. However, if the VDisk contains data, do not shrink the size of the disk.

Attention:

1. The SAN Volume Controller arbitrarily reduces the capacity of the VDisk by removing one or more extents from those that are allocated to the VDisk. You cannot control which extents are removed so you cannot guarantee that it is unused space that is removed.
2. If the VDisk contains data that is being used, *do not attempt under any circumstances to shrink a VDisk without first backing up your data.*
3. For performance reasons, some operating systems or file systems use the outer edge of the disk.

You can use the **shrinkvdisksize** command to shrink the physical capacity that is allocated to the particular VDisk by the specified amount. You can also shrink the virtual capacity of a thin-provisioned volume without altering the physical capacity assigned to the VDisk.

For more information about the command parameters, see the *IBM System Storage SAN Volume Controller Command-Line Interface User's Guide*.

Perform the following steps to shrink a VDisk:

1. Validate that the VDisk is not mapped to any host objects. If the VDisk is mapped, data is displayed.
2. You can determine the exact capacity of the source or master VDisk. Issue the following command:

```
svcinfo lsvdisk -bytes vdiskname
```

3. Shrink the VDisk by the required amount. Issue the following command:

```
svctask shrinkvdisksize -size capacitytoshrinkby -unit  
unitsforreduction vdiskname/ID
```

Migrating extents using the CLI

To improve performance, you can migrate extents using the command-line interface (CLI).

- The SAN Volume Controller provides various data migration features. These can be used to move the placement of data both *within* MDisk groups (storage pools) and *between* MDisk groups (storage pools). These features can be used concurrently with I/O operations. You can use either of the following ways to migrate data:
- | 1. Migrating data (extents) from one MDisk to another (within the same MDisk group (storage pool)). This can be used to remove highly utilized MDisks.
 - | 2. Migrating VDIsks (volumes) from one MDisk group (storage pool) to another. This can be used to remove highly utilized MDisk groups (storage pools). For example, you can reduce the utilization of a group of MDisks.

- | Migration commands fail if the target or source VDisk (volume) is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

You can determine the usage of particular MDisks by gathering input/output (I/O) statistics about nodes, MDisks, and VDIsks (volumes). After you have gathered this data, you can analyze it to determine which MDisks are highly utilized. The procedure then takes you through querying and migrating extents to elsewhere in the same MDisk group (storage pool). This procedure can only be performed using the command-line tools.

If performance monitoring tools, such as IBM Tivoli Storage Productivity Center, indicate that a managed disk in the pool is being overutilized, you can migrate some of the data onto other MDisks within the same MDisk group (storage pool).

- | 1. Determine the number of extents that are in use by each VDisk (volume) for the given MDisk by issuing the following CLI command:

```
svcinfolsmdiskextent mdiskname
```

This command returns the number of extents that each VDisk (volume) is using on the given MDisk. You should pick some of these to migrate elsewhere in the group.

- | 2. Determine the other MDisks that reside in the same MDisk group (storage pool).
 - | a. To determine the MDisk group (storage pool) that the MDisk belongs to, issue the following CLI command:

```
svcinfolsmdisk mdiskname | ID
```

- | b. List the MDisks in the group by issuing the following CLI command:

```
svcinfolsmdisk -filtervalue mdisk_grp_name=mdiskgrpname
```

- | 3. Select one of these MDisks as the target MDisk for the extents. You can determine how many free extents exist on an MDisk by issuing the following CLI command:

```
svcinfolsfreeextents mdiskname
```

You can issue the **svcinfolsmdiskextent *newmdiskname*** command for each of the target MDisks to ensure that you are not just moving the over-utilization to another MDisk. Check that the VDisk (volume) that owns the set of extents to be moved does not already own a large set of extents on the target MDisk.

- | 4. For each set of extents, issue the following CLI command to move them to another MDisk:

```
svctask migrateexts -source mdiskname | ID -exts num_extents  
-target newmdiskname | ID -threads 4 vdiskid
```

where *num_extents* is the number of extents on the *vdiskid*. The *newmdiskname* | *ID* value is the name or ID of the MDisk to migrate this set of extents to.

Note: The number of threads indicates the priority of the migration processing, where 1 is the lowest priority and 4 is the highest priority.

- | 5. Repeat the previous steps for each set of extents that you are moving.
- | 6. You can check the progress of the migration by issuing the following CLI command:

```
svcinfolsmigrate
```

Migrating VDisks (volumes) between MDisk groups using the CLI

- | You can migrate virtual disks (VDisks) between managed disk (MDisk) groups (storage pools) using the command-line interface (CLI).

You can determine the usage of particular MDisks by gathering input/output (I/O) statistics about nodes, MDisks, and VDisks (volumes). After you have gathered this data, you can analyze it to determine which VDisks or MDisks are hot. You can then migrate VDisks from one MDisk group to another.

Perform the following step to gather statistics about MDisks and VDisks:

1. Use secure copy (**scp** command) to retrieve the dump files for analyzing. For example, issue the following:

```
scp clusterip:/dumps/iostats/v_*
```

This copies all the VDisk statistics files to the AIX host in the current directory.

2. Analyze the dumps to determine which VDisks are hot. It might be helpful to also determine which MDisks are being used heavily as you can spread the data that they contain more evenly across all the MDisks in the group by migrating the extents.

After you analyze the I/O statistics data, you can determine which VDisks are hot. You also need to determine the MDisk group that you want to move this VDisk to. Either create a new MDisk group or determine an existing group that is not yet overly used. To do this, check the I/O statistics files that you generated and then ensure that the MDisks or VDisks in the target MDisk group are used less than those in the source group.

You can use data migration or VDisk mirroring to migrate data between MDisk groups. Data migration uses the command **svctask migratevdisk**. VDisk mirroring uses the commands **svctask addvdiskcopy** and **svctask rmvdiskcopy**.

When you issue the **svctask migratevdisk** command, a check is made to ensure that the destination of the migration has enough free extents to satisfy the command. If it does, the command proceeds. The command takes some time to complete.

Notes:

- You cannot use the SAN Volume Controller data migration function to move a VDisk between MDisk groups that have different extent sizes.
- Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

When you use data migration, it is possible for the free destination extents to be consumed by another process; for example, if a new VDisk is created in the destination MDisk group or if more migration commands are started. In this scenario, after all the destination extents are allocated, the migration commands suspend and an error is logged (error ID 020005). To recover from this situation, use either of the following methods:

- Add additional MDisks to the target MDisk group. This provides additional extents in the group and allows the migrations to be restarted. You must mark the error as fixed before you reattempt the migration.
- Migrate one or more VDisks that are already created from the MDisk group to another group. This frees up extents in the group and allows the original migrations to be restarted.

Perform the following steps to use the **svctask migratevdisk** command to migrate VDisks between MDisk groups:

1. After you determine the VDisk that you want to migrate and the new MDisk group you want to migrate it to, issue the following CLI command:

```
svctask migratevdisk -vdisk vdiskname/ID -mdiskgrp
newmdiskgrname/ID -threads 4
```

2. You can check the progress of the migration by issuing the following CLI command:

```
svcinfolsmigrate
```

When you use data migration, the VDisk goes offline if either MDisk group fails. VDisk mirroring can be used to minimize the impact to the VDisk because the VDisk goes offline only if the source MDisk group fails.

Perform the following steps to use VDisk mirroring to migrate VDIs between MDisk groups:

1. After you determine the VDisk that you want to migrate and the new MDisk group that you want to migrate it to, issue the following command:

```
svctask addvdiskcopy -mdiskgrp newmdiskgrname/ID vdiskname/ID
```

2. The copy ID of the new copy is returned. The copies now synchronize such that the data is stored in both MDisk groups. You can check the progress of the synchronization by issuing the following command:

```
svcinfolsvdisksyncprogress
```

3. After the synchronization is complete, remove the copy from the original I/O group to free up extents and decrease the utilization of the MDisk group. To remove the original copy, issue the following command:

```
svctask rmvdiskcopy -copy original copy id vdiskname/ID
```

Migrating a VDisk between I/O groups using the CLI

Ensure that you are familiar with migrating a virtual disk (VDisk) between I/O groups.

Attention:

- These migration tasks are disruptive. The cached data that is held within the cluster must first be written to disk before the allocation of the VDisk (volume) can be changed.
- Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

Modifying the I/O group that services the VDisk cannot be done concurrently with I/O operations. It also requires a rescan at the host level to ensure that the multipathing driver is notified that the allocation of the preferred node has changed and the ports by which the VDisk is accessed has changed. This should only be done in the situation where one pair of nodes has become over utilized.

Perform the following steps to migrate a VDisk between I/O groups:

1. Synchronize all file systems that are mounted on the given VDisk.
2. Stop all I/O operations to the VDisk.
3. Issue the following CLI command to migrate the VDisk into a new I/O group:

```
svctask chvdisk -iogrp iogrp_name_or_id -node preferred_node vdisk
```

where *iogrp_name_or_id* is the name or ID of the I/O group that you want to migrate the VDisk to, *preferred_node* is the name of the node that you want to move the VDisk to, and *vdisk* is the name of the VDisk that you want to migrate.

4. Resynchronize the VDisk to host mapping. See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* or the documentation that is provided with your multipathing driver for more information.
5. Restart the I/O operations to the VDisk.

Creating an image mode VDisk (volume) using the CLI

You can use the command-line interface (CLI) to import storage that contains existing data and continue to use this storage. You can also use the advanced functions, such as Copy Services, data migration, and the cache. These disks are known as image mode virtual disks (VDisks).

Make sure you are aware of the following before you create image mode VDisks (volumes):

1. Unmanaged-mode managed disks (MDisks) that contain existing data cannot be differentiated from unmanaged-mode MDisks that are blank. Therefore, it is vital that you control the introduction of these MDisks to the cluster by adding these disks one at a time. For example, map a single LUN from your RAID controller to the cluster and refresh the view of MDisks. The newly detected MDisk is displayed.
2. *Do not* manually add an unmanaged-mode MDisk that contains existing data to an MDisk group (storage pool). If you do, the data is lost. When you use the command to convert an image mode VDisk from an unmanaged-mode disk, you will select the MDisk group where it should be added.

See the following Web site for more information:

Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

Perform the following steps to create an image mode VDisk:

1. Stop all I/O operations from the hosts. Unmap the logical disks that contain the data from the hosts.
2. Create one or more MDisk groups.
3. Map a single RAID array or logical unit from your RAID controller to the cluster. You can do this through a switch zoning or a RAID controller based on your host mappings. The array or logical unit appears as an unmanaged-mode MDisk to the SAN Volume Controller.
4. Issue the `svcinfolsmdisk` command to list the unmanaged-mode MDisks.

If the new unmanaged-mode MDisk is not listed, you can perform a fabric-level discovery. Issue the `svctask detectmdisk` command to scan the fibre-channel network for the unmanaged-mode MDisks.

Note: The `svctask detectmdisk` command also rebalances MDisk access across the available controller device ports.

5. Convert the unmanaged-mode MDisk to an image mode virtual disk.

Note: If the VDisk that you are converting maps to a solid-state drive (SSD), the data that is stored on the VDisk is not protected against SSD failures or node failures. To avoid data loss, add a VDisk copy that maps to an SSD on another node.

Issue the `svctask mkvdisk` command to create an image mode virtual disk object.

6. Map the new VDisk to the hosts that were previously using the data that the MDisk now contains. You can use the `svctask mkvdiskhostmap` command to create a new mapping between a VDisk and a host. This makes the image mode VDisk accessible for I/O operations to the host.

After the VDisk is mapped to a host object, the VDisk is detected as a disk drive with which the host can perform I/O operations.

If you want to virtualize the storage on an image mode VDisk, you can transform it into a striped VDisk. Migrate the data on the image mode VDisk to managed-mode disks in another MDisk group. Issue the `svctask migratevdisk` command to migrate an entire image mode VDisk from one MDisk group to another MDisk group.

Migrating data to an image mode virtual disk using the CLI

You can use the command-line interface (CLI) to migrate data to an image mode virtual disk (VDisk).

The **svctask migratetoimage** CLI command allows you to migrate the data from an existing VDisk (volume) onto a different managed disk (MDisk).

When the **svctask migratetoimage** CLI command is issued, it migrates the data of the user specified source VDisk onto the specified target MDisk. When the command completes, the VDisk is classified as an image mode VDisk.

Note: Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

The MDisk specified as the target must be in an unmanaged state at the time the command is run. Issuing this command results in the inclusion of the MDisk into the user specified MDisk group.

Issue the following CLI command to migrate data to an image mode VDisk:

```
svctask migratetoimage -vdisk vdiskname/ID  
-mdisk newmdiskname/ID -mdiskgrp newmdiskgrpname/ID  
-threads 4
```

where *vdiskname/ID* is the name or ID of the VDisk, *newmdiskname/ID* is the name or ID of the new MDisk, and *newmdiskgrpname/ID* is the name or ID of the new MDisk group (storage pool).

Deleting a node from a cluster using the CLI

You can use the command-line interface (CLI) to remove a node from a cluster.

After the node is deleted, the other node in the I/O group enters write-through mode until another node is added back into the I/O group.

By default, the **rmnode** command flushes the cache on the specified node before taking the node offline. When operating in a degraded state, the SAN Volume Controller ensures that data loss does not occur as a result of deleting the only node with the cache data.

Attention:

- If you are removing a single node and the remaining node in the I/O group is online, the data can be exposed to a single point of failure if the remaining node fails.
- If both nodes in the I/O group are online and the VDisks (volumes) are already degraded before deleting the node, redundancy to the volumes is already degraded. Removing a node might result in loss of access to data, and data loss might occur if the **force** option is used.
- Removing the last node in the cluster destroys the cluster. Before you delete the last node in the cluster, ensure that you want to destroy the cluster.
- When you delete a node, you remove all redundancy from the I/O group. As a result, new or existing failures can cause I/O errors on the hosts. The following failures can occur:
 - Host configuration errors
 - Zoning errors
 - Multipathing software configuration errors
- If you are deleting the last node in an I/O group and there are volumes assigned to the I/O group, you cannot delete the node from the cluster if the node is online. You must back up or migrate all data that you want to save before you delete the node. If the node is offline, you can delete the node.
- To take the specified node offline immediately without flushing the cache or ensuring that data loss does not occur, run the **rmnode** command with the **force** parameter. The **force** parameter forces continuation of the command even though any node-dependent volumes will be taken offline. Use the **force** parameter with caution; access to data on node-dependent volumes will be lost.

Perform the following steps to delete a node:

1. If you are deleting the last node in an I/O group, determine the volumes that are still assigned to this I/O group:
 - a. Issue the following CLI command to request a filtered view of the volumes:

```
svcinfolsvdisk -filtervalue IO_group_name=name
```

Where *name* is the name of the I/O group.
 - b. Issue the following CLI command to list the hosts that this volume is mapped to:

```
svcinfolsvdiskhostmap vdiskname/id
```

Where *vdiskname/id* is the name or ID of the volume.

Note: If volumes are assigned to this I/O group that contain data that you want to continue to access, back up the data or migrate the volumes to a different (online) I/O group.

2. If this node is *not* the last node in the cluster, turn off the power to the node that you intend to remove. This step ensures that the multipathing device driver, such as the subsystem device driver (SDD), does not rediscover the paths that are manually removed before you issue the delete node request.

Attention:

- a. If you are removing the configuration node, the **rmnode** command causes the configuration node to move to a different node within the cluster. This process might take a short time, typically less than a minute. The cluster IP address remains unchanged, but any SSH client attached to the configuration node might must reestablish a connection.
 - b. If you turn on the power to the node that has been removed and it is still connected to the same fabric or zone, it attempts to rejoin the cluster. The cluster causes the node to remove itself from the cluster and the node becomes a candidate for addition to this cluster or another cluster.
 - c. If you are adding this node into the cluster, ensure that you add it to the same I/O group that it was previously a member of. Failure to do so can result in data corruption.
 - d. In a service situation, a node should normally be added back into a cluster using the original node name. As long as the partner node in the I/O group has not been deleted too, this is the default name used if **-name** is not specified.
3. Before you delete the node, update the multipathing device driver configuration on the host to remove all device identifiers that are presented by the volumes that you intend to remove. If you are using the subsystem device driver, the device identifiers are referred to as virtual paths (vpaths).

Attention: Failure to perform this step can result in data corruption.

See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.

4. Issue the following CLI command to delete a node from the cluster:

Attention: Before you delete the node: The **rmnode** command checks for node-dependent volumes, which are not mirrored at the time that the command is run. If any node-dependent volumes are found, the command stops and returns a message. To continue removing the node despite the potential loss of data, run the **rmnode** command with the **force** parameter. Alternatively, follow these steps before you remove the node to ensure that all volumes are mirrored:

1. a. Run the **lsindependentvdisks** command.
- b. For each node-dependent volume that is returned, run the **lsvdisk** command.
- c. Ensure that each volume returns in-sync status.

```
svctask rmnode node_name_or_id
```

Where *node_name_or_id* is the name or ID of the node.

Performing the cluster maintenance procedure using the CLI

You can use the command-line interface (CLI) to perform the cluster maintenance procedure.

Perform the following steps for cluster maintenance:

1. Issue the **svctask finderr** command to analyze the error log for the highest severity of unfixed errors. This command scans the error log for any unfixed errors. Given a priority ordering defined within the code, the highest priority of unfixed errors is returned.
2. Issue the **svctask dumperrlog** command to dump the contents of the error log to a text file.
3. Locate and fix the error.
4. Issue the **svctask clearerrlog** command to clear all entries from the error log, including status events and any unfixed errors. Only issue this command when you have either rebuilt the cluster or have fixed a major problem that has caused many entries in the error log that you do not want to fix individually.

Note: Clearing the error log does not fix the errors.

5. Issue the **svctask cherrstate** command to toggle the state of an error between unfixed and fixed.

Modifying the cluster IP addresses using the CLI

| You can use the command-line interface (CLI) to change the IP addresses that are associated with a cluster.
|

Attention: When you specify a new IP address for a cluster, the existing communication with the cluster is broken. You must reconnect to the cluster with the new IP address.

Perform the following steps to change the cluster IP address:

1. Issue the `svcinfolclusterip` command to list the current IP addresses that are used by the cluster.
2. Record the current IP addresses for future reference.
3. To change an IPv4 cluster IP address, issue the following command:

```
svctask chclusterip -clusterip cluster_ip_address -port cluster_port
```

where *cluster_ip_address* is the new IP address for the cluster and *cluster_port* specifies which port (1 or 2) to apply changes to.
4. To change an IPv4 cluster IP address to an IPv6 cluster IP address, issue the following command:

```
svctask chclusterip -clusterip_6 cluster_ip_address -port cluster_port
```

where *cluster_ip_address* is the new IPv6 address for the cluster and *cluster_port* specifies which port (1 or 2) to apply changes to.
5. To change an IPv4 default gateway IP address, issue the following command:

```
svctask chclusterip -gw cluster_gateway_address -port cluster_port
```

where *cluster_gateway_address* is the new gateway address for the cluster and *cluster_port* specifies which port (1 or 2) to apply changes to.
6. To change an IPv6 default gateway address, issue the following command:

```
svctask chclusterip -gw_6 cluster_gateway_address -port cluster_port
```

where *cluster_gateway_address* is the new gateway address for the cluster and *cluster_port* specifies which port (1 or 2) to apply changes to.
7. Issue the following command to change an IPv4 cluster subnet mask:

```
svctask chclusterip -mask cluster_subnet_mask -port cluster_port
```

where *cluster_subnet_mask* is the new subnet mask for the cluster and *cluster_port* specifies which port (1 or 2) to apply changes to.
8. For IPv6 addresses, you can issue the following command to set the prefix for the cluster:

```
svctask chclusterip -prefix_6 -port cluster_port
```

where *cluster_port* specifies which port (1 or 2) to apply changes to.
9. Optionally, if you want to delete all of the IPv4 addresses in the cluster after you have changed all addresses to IPv6, issue the following command:

```
svctask chcluster -noip
```
10. Optionally, if you want to delete all of the IPv6 addresses in the cluster after you have changed all addresses to IPv4, issue the following command:

```
svctask chcluster -noip_6
```
11. The IP routing table provides details of the gateway that is used for IP traffic to a range of IP addresses for each Ethernet port. This information can be used to diagnose configuration node accessibility problems. To display the IP routing table, enter the following CLI command:

```
svcinfolroute
```
12. The ping command can be used to diagnose IP configuration problems by checking whether a given IP address is accessible from the configuration node. The command can be useful for diagnosing problems where the configuration node cannot be reached from a specific management server. For example, enter the following CLI command:

```
svctask ping ipv4_address | ipv6_address
```

where *ipv4_address* | *ipv6_address* is either the IPv4 address or the IPv6 address.

Changing the cluster gateway address using the CLI

- | You can use the command-line interface (CLI) to change the gateway address for a cluster.

Perform the following steps to change the cluster gateway address:

1. Issue the **svcinfolclusterip** command to list the current gateway address of the cluster.
2. Record the current gateway address for future reference.
3. Issue the following command to change an IPv4 cluster gateway address:

```
svctask chclusterip -gw cluster_gateway_address -port cluster_port
```

where *cluster_gateway_address* is the new gateway address for the cluster. The **port** parameter specifies which port (1 or 2) to apply changes to.

4. Issue the following command to change an IPv6 cluster gateway address:

```
svctask chclusterip -gw_6 cluster_gateway_address -port cluster_port
```

where *cluster_gateway_address* is the new gateway address for the cluster. The **port** parameter specifies which port (1 or 2) to apply changes to.

Changing the relationship bandwidth for a cluster using the CLI

- | You can use the command-line interface (CLI) to change the relationship bandwidth for a cluster.

The relationship bandwidth limit controls the maximum rate at which any one remote-copy relationship can synchronize. The overall limit is controlled by the **bandwidth** parameter of each cluster partnership. The default value for the relationship bandwidth limit is 25 megabytes per second (MBps), but you can change this by following these steps:

1. Issue the **svcinfolcluster** command to list the current relationship bandwidth limit of the cluster. For example:

```
svcinfolcluster cluster_id_or_cluster_name
```

Where *cluster_id_or_cluster_name* is the ID or name of the cluster.

2. For future reference, record the current relationship bandwidth limit that is displayed. For example:
`relationship_bandwidth_limit 25`
3. To change the relationship bandwidth limit of the cluster, issue the following command:

```
svctask chcluster -relationshipbandwidthlimit  
cluster_relationship_bandwidth_limit
```

Where *cluster_relationship_bandwidth_limit* is the new limit for the cluster. Issue the command on both clusters in a relationship.

Configuring the cluster for iSCSI using the CLI

You need to complete several tasks to configure the cluster to work with iSCSI-attached hosts. The tasks include general tasks on the host system before you configure cluster on the SAN Volume Controller.

Before completing any iSCSI-configuration tasks on the cluster, it is important that you complete all the iSCSI-related configuration on the host machine. Because the SAN Volume Controller supports a variety of host machines, consult the documentation for specific instructions and requirements for a particular host. For a list of supported hosts, see this website:

Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

To configure a cluster for iSCSI, follow these general tasks on the host system:

1. Select a software-based iSCSI initiator, such as Microsoft Windows iSCSI Software Initiator and verify the iSCSI driver installation.
2. If required, install and configure a multipathing driver for the host system.

In addition, determine a naming convention for iSCSI names, such as iSCSI qualified names (IQNs) for your cluster. Hosts use iSCSI names to connect to the node. Each node, for example, has a unique IQN, and the cluster name and node name are used as part of that IQN. Each node, for example, has a unique IQN, and the cluster name and node name are used as part of that IQN.

Port IP addresses are the IP addresses that are used by iSCSI-attached hosts to perform I/O.

1. To configure a new port IP address to a specified Ethernet port of a node with an IPv4 address, enter the following command-line interface (CLI) command:

```
svctask cfgportip -node node_name | node_id -ip ipv4addr
-gw ipv4gw -mask subnet_mask -failover port_id
```

where *node_name* | *node_id* specifies the name or ID of the node that is being configured, *ipv4addr* is the IPv4 address for the Ethernet port, *ipv4gw* is the IPv4 gateway IP address, *subnet_mask* is the IPv4 subnet mask, and *port_id* specifies the Ethernet port ID (1 or 2). To view a list of ports, use the **svcinfolspportip** command.

The optional **-failover** parameter specifies that the data is failover data, which is data that is related to the partner node. If the node that is specified is the only online node in the I/O group, the address is configured and presented by this node. When another node in the I/O group comes online, the failover address is presented by that node. If two nodes in the I/O group are online when the command is issued, the address is presented by the other node to that specified.

2. To configure a new port IP address that belongs to a partner node with an IPv6 address in the I/O group, enter the following CLI command:

```
svctask cfgportip -node node_name | node_id -ip_6 ipv6addr
-gw_6 ipv6gw -prefix_6 prefix -failover port_id
```

where *node_name* | *node_id* specifies the name or ID of the node that is being configured, *ipv6addr* is the IPv6 address for the iSCSI port, *ipv6gw* if the gateway address for the given IP address, *prefix* is the IPv6 prefix for the gateway, and *port_id* specifies the Ethernet port ID (1 or 2). To view a list of ports, use the **svcinfolspportip** command. If the partner node is offline, the address is configured and presented by this node. When another node comes online in the I/O group, the failover address is presented by that node.

The optional **-failover** parameter specifies that the data is failover data, which is data that is related to the partner node. If the node that is specified is the only online node in the I/O group, the address is configured and presented by this node. When another node in the I/O group comes online, the failover address is presented by that node. If two nodes in the I/O group are online when the command is issued, the address is presented by the other node to that specified.

3. To remove an iSCSI IP address from a node Ethernet port, enter either of these CLI commands. The following command deletes an IPv4 configuration for the specified iSCSI Ethernet port:

```
svctask rmpportip -failover
-node node_name | node_id port_id
```

where *node_name* | *node_id* specifies the name or ID of the node with the Ethernet port that the IP address is being removed from and *port_id* specifies the Ethernet port ID. To list the valid values for the Ethernet port, enter the **svcinfolspportip** command. The optional **-failover** parameter indicates that the specified data is failover data.

The following command deletes an IPv6 configuration for the specified iSCSI Ethernet port:

```
svctask rmpportip -ip_6 -failover
-node node_name | node_id port_id
```

where **-ip_6** indicates that this command will remove an IPv6 configuration, *node_name* | *node_id* specifies the name or ID of the node with the Ethernet port that the IP address is being removed

from, and *port_id* specifies the Ethernet port ID. To list the valid values for the Ethernet port, enter the **svcinfolspportip** command. The optional **-failover** parameter indicates that the specified data is failover data.

After you configure your IP addresses, you can optionally create iSCSI aliases.

Configuring or modifying an iSCSI alias using the CLI

You can use the command-line interface (CLI) to optionally create or change the iSCSI alias for the selected node. An iSCSI alias is a user-assigned name that identifies the SAN Volume Controller node to the iSCSI-attached host. It is possible to change the iSCSI alias of a node even if the node is offline.

To configure or modify an iSCSI alias, follow these steps:

1. To configure a new port IP address to a specified Ethernet port of a node, enter the following CLI command:

```
svctask chnode -iscsialias alias node_name | node_id
```

where *alias node_name* | *node_id* specifies the name or ID of the node.

2. To specify that the name or iSCSI alias that is being set is the name or alias of the partner node in the I/O group, enter the following CLI command. When there is no partner node, the values set are applied to the partner node when it is added to the cluster. If this parameter is used when there is a partner node, the name or alias of that node changes

```
svctask chnode -iscsialias alias -failover node_name | node_id
```

where *alias* specifies the iSCSI name of the node and *node_name* | *node_id* specifies the node to be modified.

After you create iSCSI aliases, you can optionally configure the address for the Internet Storage Name Service (iSNS) server for the cluster.

Configuring the iSNS server address using the CLI

If you are using iSCSI-attached hosts with the SAN Volume Controller cluster, you can the command-line interface (CLI) to optionally configure the address for the Internet Storage Name Service (iSNS) server for the cluster. Host systems use the iSNS server to manage iSCSI targets and for iSCSI discovery.

1. To specify an IPv4 address for the iSCSI storage name service (SNS), enter the following CLI command:

```
svctask chcluster -isnsip sns_server_address
```

where *sns_server_address* is the IP address of the iSCSI storage name service in IPv4 format.

2. To specify an IPv6 address for the iSCSI storage name service (SNS), enter the following CLI command:

```
svctask chcluster -isnsip_6 ipv6_sns_server_address
```

where *ipv6_sns_server_address* is the IP address of the iSCSI storage name service in IPv6 format.

After you configure the iSNS server address for the cluster, you can configure cluster iSCSI authentication.

Note: To help in problem determination, this step can be delayed until after the first one or two hosts have been configured and their connectivity has been tested without authentication configured.

Configuring cluster iSCSI authentication using the CLI

You can use the command-line interface (CLI) to configure the Challenge-Handshake Authentication Protocol (CHAP) to authenticate the SAN Volume Controller cluster to the iSCSI-attached hosts. After the CHAP is set for the cluster, all attached hosts must be configured to authenticate this way. To help in

problem determination, this step can be delayed until after the first one or two hosts have been configured and their connectivity has been tested without authentication configured.

To configure authentication between the SAN Volume Controller cluster and the iSCSI-attached hosts, follow these steps:

1. To set the authentication method for the iSCSI communications of the cluster, enter the following CLI command:

```
svctask chcluster -iscsiauthmethod chap -chapsecret chap_secret
```

where *chap* sets the authentication method for the iSCSI communications of the cluster and *chap_secret* sets the CHAP secret to be used to authenticate the cluster via iSCSI. This parameter is required if the **iscsiauthmethod chap** parameter is specified. The specified CHAP secret cannot begin or end with a space.

2. To clear any previously set CHAP secret for iSCSI authentication, enter the following CLI command:
svctask chcluster -nochapsecret

The **nochapsecret** parameter is not allowed if the **chapsecret** parameter is specified.

3. The **lsiscsiauth** command lists the Challenge Handshake Authentication Protocol (CHAP) secret that is configured for authenticating an entity to the SAN Volume Controller cluster. The command also displays the configured iSCSI authentication method. For example, enter the following CLI command:

```
svcinfolscsiauth
```

After you configure the CHAP secret for the SAN Volume Controller cluster, ensure that the cluster CHAP secret is added to each iSCSI-attached host. On all iSCSI-attached hosts, specify a CHAP secret that the hosts use to authenticate to the SAN Volume Controller cluster.

Configuring remote authentication service using CLI

You can use the command-line interface (CLI) to configure the SAN Volume Controller to use remote authentication.

| Remote authentication allows users of SAN management applications, such as IBM Tivoli Storage
| Productivity Center, to authenticate to the cluster using the authentication service provided by the SAN
| management application. To use the SAN Volume Controller with a remote authentication service, follow
| these steps:

1. Configure the cluster with the location of the remote authentication server.

To change settings, issue the **svctask chauthservice** command. To view settings, issue the **svcinfolsccluster** command.

You can use either an http or https connection to the server. If you use an http option, the user and password information is transmitted in clear text over the IP network.

2. Configure user groups on the cluster by matching those that are used by the authentication service.

For each group of interest known to the authentication service, a SAN Volume Controller user group must be created with the same name and with the remote setting enabled. If, for example, members of a group called *sysadmins* require the SAN Volume Controller Administrator (admin) role, issue the following command:

```
svctask mkusergrp -name sysadmins -remote -role Administrator
```

If none of the groups for a user match any of the SAN Volume Controller user groups, the user is not permitted to access the cluster.

3. Configure users who do not require Secure Shell (SSH) access.

SAN Volume Controller users who are to use the remote authentication service and do not require SSH access should be deleted from the system. The superuser cannot be deleted and cannot use the remote authentication service.

4. Configure users who require SSH access.

All SAN Volume Controller users who are to use the remote authentication service and require SSH access must have their remote settings enabled and the same password set both on the cluster and on the authentication service.

The remote setting instructs SAN Volume Controller to check the authentication service for group information for determining the role of the user.

5. Configure the system time.

The current time of both the SAN Volume Controller cluster and the system that is running the remote authentication service must match. The easiest way to do this is to use the same Network Time Protocol (NTP) server for both.

Attention: Failure to follow this step could result in either poor interactive performance of the SAN Volume Controller user interface or in incorrect user-role assignments.

Creating and working with user groups using the CLI

| You can use the command-line interface (CLI) to create and work with users and user groups. User
| groups organize users of a cluster by role.

| Roles apply to both local and remote users on the cluster and are based on the user group to which the
| user belongs. A local user can only belong to a single group; therefore, the role of a local user is defined
| by the single group that the user belongs to. Remote users can belong to one or more groups; therefore,
| the roles of remote users are assigned according to the groups that the remote user belongs to.

To create and work with user groups, follow these steps:

1. Issue the **svctask mkusergrp** CLI command to create a new user group. For example:

```
svctask mkusergrp -name group_name -role role_name -remote
```

where *group_name* specifies the name of the user group and *role_name* specifies the role that is associated with any users that belong to this group. The **remote** parameter specifies that the group is visible to the remote authentication service.

| The command returns the ID of the user group that was created. To create user groups in the
| management GUI, select **User Management** → **Users**. From the **Global Actions** menu, select **New User**
| **Group**.

2. Issue the **svctask chusergrp** CLI command to change attributes of an existing user group. For example:

```
svctask chusergrp -role role_name -remote yes | no group_id_or_name
```

where *role_name* specifies the role that is associated with any users that belong to this group and *group_id_or_name* specifies the group to be changed. The **remote** parameter specifies whether the group is visible to the authentication server.

| To change a user group in the management GUI, select **User Management** → **Users**. Select a user
| group and select **Properties** from the **Actions** menu.

3. Issue the **svctask rmusergrp** CLI command to delete a user group: For example:

```
svctask rmusergrp -force group_id_or_name
```

where *group_id_or_name* specifies the group to delete. The **force** parameter specifies to delete the group even if there are users in the user group. All users that were assigned to this group are assigned to the Monitor group.

| To delete a user group in the management GUI, select **User Management** → **Users**. Select a user group
| and select **Delete** from the **Actions** menu.

4. Issue the **svctask lsusergrp** CLI command to display the user groups that have been created on the cluster. For example:

```
svcinfolusergrp usergrp_id_or_name
```

where *group_id_or_name* specifies the user group to view. If you do not specify a user group ID or name, all user groups on the cluster are displayed.

Creating and working with users using the CLI

You can use the command-line interface (CLI) to create and work with users.

| You can create two categories of users that access the cluster. These types are based on how the users are
| authenticated to the cluster. Local users must provide either a password, a Secure Shell (SSH) key, or
| both. Local users are authenticated through the authentication methods that are located on the SAN
| Volume Controller cluster. If the local user needs access to the management GUI, a password is needed
| for the user. If the user requires access to the command-line interface (CLI), a valid SSH key file is
| necessary. If a user is working with both interfaces, both a password and SSH key are required. Local
| users must be part of a user group that is defined on the cluster. User groups define roles that authorize
| the users within that group to a specific set of operations on the cluster.

| A remote user is authenticated on a remote service that is usually provided by a SAN management
| application, such as IBM Tivoli Storage Productivity Center, and does not need local authentication
| methods. For a remote user, both a password and SSH key are required to use the command-line
| interface. Remote users only need local credentials to access to the management GUI if the remote service
| is down. Remote users have their groups defined by the remote authentication service.

To create and work with users, follow these steps:

1. Issue the **svctask mkuser** CLI command to create either a local user or a remote user to access the SAN Volume Controller. For example:

```
svctask mkuser -name user_name -remote
```

where *user_name* specifies the name of the user. The **remote** parameter specifies that the user authenticates to the remote authentication service.

```
svctask mkuser -name user_name -usergrp group_name_or_id
```

where *user_name* specifies the name of the user and *group_name_or_id* specifies the name or ID of the user group with which the local user is to be associated. The **usergrp** parameter specifies that the user authenticates to the cluster using cluster authentication methods.

2. Issue the **svctask chuser** CLI command to change the attributes of an existing user. For example:

```
svctask chuser -usergrp group_id_or_name user_id_or_name
```

where the *group_id_or_name* specifies the new group for the user and *user_id_or_name* specifies the user to be changed.

3. Issue the **svctask chcurrentuser** CLI command to change the attributes of the current user. For example:

```
svctask chcurrentuser -nokey
```

where the **nokey** parameter specifies that the SSH key of the user is to be deleted.

4. Issue the **svctask rmuser** CLI command to delete a user: For example:

```
svctask rmuser user_id_or_name
```

where *user_id_or_name* specifies the user to be removed.

5. Issue the **svcinfoluser** CLI command to display a list of users that have been created on the cluster. For example:

```
svcinfoluser user_id_or_name
```

where *user_id_or_name* specifies the ID or name of the user view. If you do not specify an ID or name, the concise view is displayed. If you do not specify a user ID or name, all users on the cluster are displayed.

6. Issue the **svcinfo lscurrentuser** CLI command to display the name and role of the logged-in user. For example:

```
svcinfo lscurrentuser
```

The name and the role of the user are displayed.

Setting up SNMP notifications using the CLI

- | You can set up event notifications using the command-line interface (CLI).

The notification settings apply to the entire cluster. You can specify the types of events that cause the cluster to send a notification. The cluster sends a Simple Network Management Protocol (SNMP) notification. The SNMP setting represents the type of notification.

- | SNMP is the standard protocol for managing networks and exchanging messages. SNMP enables the SAN Volume Controller to send external messages that notify personnel about an event. You can use an SNMP manager to view the messages that the SNMP agent sends.

The possible types of event notifications are error, warning, and information. Event notifications are reported to the SNMP destinations of your choice. To specify an SNMP destination, you *must* provide a valid IP address and SNMP community string.

Note: A valid community string can contain up to 60 letters or digits (most characters). A maximum of six SNMP destinations can be specified.

In configurations that use SNMP, the SAN Volume Controller uses the notifications settings to call home if errors occur. You must specify Error and send the trap to the IBM System Storage Productivity Center or the master console if you want the SAN Volume Controller to call home when errors occur.

To configure the SNMP notification settings, use the following commands:

1. To create a new SNMP server to receive notifications, use the **svctask mksnmpserver** CLI command. For example, enter one of the following commands:

```
svctask mksnmpserver -ip 9.11.255.634
```

where *9.11.255.634* is the IP addresses for this server.

```
svctask mksnmpserver -ip 9.11.255.634 -port remoteporntnumber
```

where *9.11.255.634* is the IP addresses for this server and *remoteporntnumber* is the port number for the remote SNMP server.

2. To change the settings of an existing SNMP server, enter the **svctask chsnmpserver** command. For example:

```
svctask chsnmpserver -name newserver snmp_server_name_or_id
```

where *newserver* is the new name or ID of the server and *snmp_server_name_or_id* is the name or ID of the server to be modified.

3. To remove an existing SNMP server from the system, enter the **svctask rmsnmpserver** command. For example:

```
svctask rmsnmpserver snmp_server_name_or_id
```

where *snmp_server_name_or_id* is either the name or the ID of the SNMP server to be deleted.

4. To display either a concise list or a detailed view of the SNMP servers that are detected by the cluster, enter the **svcinfolssnmpserver** command. For example, to display a concise view, enter the following command:

```
svcinfolssnmpserver -delim :
```

To display a detailed view of an SNMP server, enter the following command:

```
svcinfolssnmpserver snmp_server_name_or_id
```

Setting up syslog notifications using the CLI

- 1 You can set up syslog event notifications using the command-line interface (CLI).

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The IP network can be either IPv4 or IPv6. The system can send syslog messages that notify personnel about an event. The system can transmit syslog messages in either expanded or concise format. You can use a syslog manager to view the syslog messages that the system sends. The system uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the management GUI or the SAN Volume Controller command-line interface to configure and modify your syslog settings.

The syslog event notification settings apply to the entire cluster. You can specify the types of events that cause the cluster to send a notification. The possible types of notifications are error, warning, or information.

To specify a syslog destination, you *must* provide a valid IP address.

Note: Servers that are configured with facility values of 0 - 3 receive syslog messages in concise format. Servers that are configured with facility values of 4 - 7 receive syslog messages in fully expanded format.

The SAN Volume Controller uses the notifications settings to call home if errors occur.

To configure and work with notification settings, use the following commands:

1. Issue the **svctask mksyslogserver** CLI command to specify the action that you want to take when a syslog error or event is logged to the error log. For example, you can issue the following CLI command to set up a syslog notification:

```
svctask mksyslogserver -ip 9.11.255.634
```

where *9.11.255.634* is the IP address of the syslog server.

2. To modify a syslog notification, issue the **svctask chsyslogserver** command. For example:

```
svctask chsyslogserver -name -facility facility_number syslog_server_name_or_id
```

where *facility number* is a facility number to identify the origin of the message to the receiving server and *syslog_server_name_or_id* is the name or ID of the server to be modified.

3. To delete a syslog notification, issue the **svctask rmsyslogserver** command. For example:

```
svctask rmsyslogserver syslog_server_name_or_id
```

4. To display either a concise list or a detailed view of syslog servers that are configured on the cluster, issue the **svcinfolssyslogserver** command. For example, to display a concise view, enter the following command:

```
svcinfolssyslogserver -delim :
```

To display a detailed view of a syslog server, enter the following command:

```
svcinfolssyslogserver snmp_server_name_or_id
```

Setting up e-mail event notifications and inventory reports using the CLI

- | You can use the command-line interface (CLI) to set up your system to send event notification and
- | inventory reports to specified recipients and the IBM Support Center

To set up, manage, and activate e-mail event and inventory notifications, complete the following steps:

1. Enable your system to use the e-mail notification function. To do this, issue the **svctask mkemailserver** CLI command. Up to six SMTP e-mail servers can be configured to provide redundant access to the external e-mail network.

The following example creates an e-mail server object. It specifies the name, IP address, and port number of the SMTP e-mail server. After you issue the command, you see a message that indicates that the e-mail server was successfully created.

```
svctask mkemailserver -ip ip_address -port port_number
```

where *ip_address* specifies the IP address of a remote e-mail server and *port_number* specifies the port number for the e-mail server.

2. Add recipients of e-mail event and inventory notifications to the e-mail event notification facility. To do this, issue the **svctask mkemailuser** CLI command. You can add up to twelve recipients, one recipient at a time.

The following example adds e-mail recipient **manager2008** and designates that **manager2008** receive e-mail error-type event notifications.

```
svctask mkemailuser -address manager2008@ibm.com  
-error on -usertype local
```

3. Set the contact information that is used by the e-mail event notification facility. To do this, issue the **svctask chemail** CLI command. If you are starting the e-mail event notification facility, the **reply**, **contact**, **primary**, and **location** parameters are required. If you are modifying contact information used by the e-mail event notification facility, at least one of the parameters must be specified.

The following example sets the contact information for the e-mail recipient **manager2008**.

```
svctask chemail -reply manager2008@ibm.com -contact manager2008  
-primary 0441234567 -location 'room 256 floor 1 IBM'
```

4. Optionally, generate a report that lists e-mail event notification settings for all e-mail recipients, or change or delete e-mail recipients.
 - To generate a report that lists the e-mail event notification settings for all e-mail recipients, an individual e-mail recipient, or a specified type of e-mail recipient (local or support), issue the **svctask lsemailuser** CLI command.
 - To change the settings that are defined for a recipient, issue the **svctask chemailuser** CLI command. You must specify the user ID or name of the e-mail recipient for whom you are modifying settings.
 - To remove a previously defined e-mail recipient, issue the **svctask rmemailuser** CLI command. You must specify the user ID or name of the e-mail recipient that you want to remove.
5. Activate the e-mail and inventory notification function. To do this, issue the **svctask startemail** CLI command. There are no parameters for this command.

Note: Inventory information is automatically reported to IBM when you activate error reporting.

6. Optionally, test the e-mail notification function to ensure that it is operating correctly and send an inventory e-mail notification.
 - To send a test e-mail notification to one or more recipients, issue the **svctask testemail** CLI command. You must either specify **all** or the user ID or user name of an e-mail recipient that you want to send a test e-mail to.

- To send an inventory e-mail notification to all recipients that are enabled to receive inventory e-mail notifications, issue the **svctask sendinventoryemail** CLI command. There are no parameters for this command.

Setting up e-mail servers using the CLI

| You can set up e-mail server objects using the command-line interface (CLI).

You can specify a server object that describes a remote Simple Mail Transfer Protocol (SMTP) e-mail server to receive event notifications from the cluster. You can specify up to six servers to receive notifications. To configure and work with e-mail servers, use the following commands:

1. Issue the **svctask mkemailserver** CLI command to create an e-mail server object that describes a remote Simple Mail Transfer Protocol (SMTP) e-mail server. For example, issue the following CLI command to set up an e-mail server:

```
svctask mkemailserver -ip ip_address
```

where *ip_address* is the IP address of a remote e-mail server. This must be a valid IPv4 or IPv6 address.

2. To change the parameters of an existing e-mail server object, issue the **svctask chemailserver** command. For example:

```
svctask chemailserver -ip ip_address email_server_name_or_id
```

where *ip_address* is the IP address of the e-mail server object and *email_server_name_or_id* is the name or ID of the server object to be changed.

3. To delete a specified e-mail server object, issue the **svctask rmemailserver** command. For example:

```
svctask rmemailserver email_server_name_or_id
```

4. To display either a concise list or a detailed view of e-mail servers that are configured on the cluster, issue the **svcinfolservermailserver** command. For example, to display a concise view, enter the following command:

```
svcinfolservermailserver -delim :
```

To display a detailed view of an e-mail server, enter the following command:

```
svcinfolservermailserver email_server_name_or_id
```

Changing cluster passwords using the CLI

You can use the command-line interface (CLI) to change the superuser and service passwords.

Passwords only affect the management GUI that accesses the cluster. To restrict access to the CLI, you must control the list of SSH client keys that are installed on the cluster.

Perform the following steps to change the superuser and service passwords:

1. Issue the following command to change the superuser password:

```
svctask chuser -password superuser_password superuser
```

Where *superuser_password* is the new superuser password that you want to use.

2. Issue the following command to change the service password:

|

```
svctask chcluster -servicepwd service_password
```

Where *service_password* is the new service password that you want to use.

Changing the locale setting using the CLI

You can use the command-line interface (CLI) to specify the locale for a SAN Volume Controller cluster. The language that you select as your locale setting is used to display command results and error messages in the CLI.

The following locales are available:

- 0 US English (default)
- 3 Japanese

Issue the **svctask setlocale** CLI command with the ID for the locale.

For example, issue the following CLI command to change the locale setting from US English to Japanese:

```
svctask setlocale 3
```

where 3 is the ID for the Japanese locale setting.

Viewing the feature log using the CLI

You can use the command-line interface (CLI) to view the feature log.

Perform the following steps to view the feature log:

1. Issue the **svcinfolfeaturedumps** command to return a list of dumps in the /dumps/feature destination directory. The feature log is maintained by the cluster. The feature log records events that are generated when license parameters are entered or when the current license settings have been breached.
2. Issue the **svcservicemodeinfo lfeaturedumps** command to return a list of the files that exist of the type specified on the given node.

Analyzing the error log using the CLI

- | You can use the command-line interface (CLI) to analyze the error log (event log).

Perform the following step to analyze the error log:

- | Issue the following CLI command to list error log entries by file type: **lseventlog**
- | This command lists the error log entries. You can filter by type; for example, **lseventlog -filtervalue object_type=mdisk** displays the error log by managed disks (MDisks).

You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. You can also request that the output is sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. Therefore, the most serious errors are displayed first in the table. For time, either the older or the latest entry can be displayed first in the output.

Shutting down a cluster using the CLI

You can use the command-line interface (CLI) to shut down a cluster.

If you want to remove all input power to a cluster (for example, the machine room power must be shutdown for maintenance), you must shut down the cluster before the power is removed. If you do not shut down the cluster before turning off input power to the uninterruptible power supply, the SAN Volume Controller nodes detect the loss of power and continue to run on battery power until all data that is held in memory is saved to the internal disk drive. This increases the time that is required to make

the cluster operational when input power is restored and severely increases the time that is required to recover from an unexpected loss of power that might occur before the uninterruptible power supply batteries have fully recharged.

When input power is restored to the uninterruptible power supply units, they start to recharge. However, the SAN Volume Controller nodes do not permit any I/O activity to be performed to the virtual disks (VDisks) until the uninterruptible power supply is charged enough to enable all the data on the SAN Volume Controller nodes to be saved in the event of an unexpected power loss. This might take as long as two hours. Shutting down the cluster prior to removing input power to the uninterruptible power supply units prevents the battery power from being drained and makes it possible for I/O activity to resume as soon as input power is restored.

Before shutting down a cluster, quiesce all I/O operations that are destined for this cluster. Failure to do so can result in failed I/O operations being reported to your host operating systems.

- | **Attention:** If you are shutting down the entire cluster, you lose access to all VDisks (volumes) that are provided by this cluster. Shutting down the cluster also shuts down all SAN Volume Controller nodes. This shutdown causes the hardened data to be dumped to the internal hard drive.

Begin the following process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks that are provided by the cluster.

1. Determine which hosts are using the VDisks that are provided by the cluster.
2. Repeat the previous step for all VDisks.

If input power is lost and subsequently restored, you must press the power button on the uninterruptible power supply units before you press the power buttons on the SAN Volume Controller nodes.

Perform the following steps to shut down a cluster:

1. Issue the following command to shut down a cluster:
`svctask stopcluster`

The following output is displayed:

```
Are you sure that you want to continue with the shut down?
```

2. Type `y` to shut down the entire cluster.

Upgrading the software automatically using the CLI

You can use the command-line interface (CLI) to install software upgrades.

This procedure is for upgrading from SAN Volume Controller version 6.1.0 or later. To upgrade from version 5.1.x or earlier, see the relevant information center or publications that are available at this website:

Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

- | **Attention:** Before you start a software upgrade, you must check for offline or degraded volumes. An offline volume can cause write data that has been modified to be pinned in the SAN Volume Controller cache. This prevents volume failover and causes a loss of I/O access during the software upgrade. If the `fast_write_state` is empty, a volume can be offline and not cause errors during the software upgrade.

Perform the following steps to upgrade the software:

1. Download, install, and run the latest version of the Software Upgrade Test Utility to verify that there are no issues with the current cluster environment. You can download the most current version of this tool at the following website:

<http://www.ibm.com/support/docview.wss?uid=ssg1S4000585>

2. Download the SAN Volume Controller code from the Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145.
 - If you want to write the SAN Volume Controller code to a CD, you must download the CD image.
 - If you do not want to write the SAN Volume Controller code to a CD, you must download the installation image.
3. Use PuTTY scp (pscp) to copy the software upgrade files to the node.
4. Ensure that the software upgrade file has been successfully copied.

Before you begin the software upgrade, you must be aware of the following:

- The installation process fails under the following conditions:
 - If the software that is installed on the remote cluster is not compatible with the new software or if there is an intercluster communication error that does not allow the software to check that the software is compatible.
 - If any node in the cluster has a hardware type that is not supported by the new software.
 - If the SAN Volume Controller software determines that one or more volumes in the cluster would be taken offline by rebooting the nodes as part of the upgrade process. You can find details about which volumes would be affected by using the **svcinfoldependentvdisks** command. If you are prepared to lose access to data during the upgrade, you can use the force flag to override this restriction.
- The software upgrade is distributed to all the nodes in the cluster using fibre-channel connections between the nodes.
- Nodes are updated one at a time.
- Nodes will run the new software concurrently with normal cluster activity.
- While the node is updated, it does not participate in I/O activity in the I/O group. As a result, all I/O activity for the volumes in the I/O group is directed to the other node in the I/O group by the host multipathing software.
- There is a 30-minute delay between node updates. The delay allows time for the host multipathing software to rediscover paths to the nodes which have been upgraded, so that there is no loss of access when another node in the I/O group is updated.
- The software update is not committed until all nodes in the cluster have been successfully updated to the new software level. If all nodes successfully restart with the new software level, the new level is committed. When the new level is committed, the cluster vital product data (VPD) is updated to reflect the new software level.
- You cannot invoke the new functions of the upgraded software until all member nodes are upgraded and the update has been committed.
- Because the software upgrade process takes some time, the installation command completes as soon as the software level is verified by the cluster. To determine when the upgrade has completed, you must either display the software level in the cluster VPD or look for the Software upgrade complete event in the error/event log. If any node fails to restart with the new software level or fails at any other time during the process, the software level is backed-off.
- During a software upgrade, the version number of each node is updated when the software has been installed and the node has been restarted. The cluster software version number is updated when the new software level is committed.
- When the software upgrade starts an entry is made in the error or event log and another entry is made when the upgrade completes or fails.

5. Issue the following CLI command to start the software upgrade process:

```
svctask applysoftware -file software_upgrade_file
```

where *software_upgrade_file* is the name of the software upgrade file. If the cluster identifies any volumes that would go offline as a result of rebooting the nodes as part of the cluster upgrade, the software upgrade does not start. An optional force parameter can be used to indicate that the upgrade continue in spite of the problem identified. Use the **svcinfoldependentvdisks** command to identify the cause for the failed upgrade. If you use the force parameter, you are prompted to confirm that you want to continue. The behavior of the force parameter has changed, and it is no longer required when applying an upgrade to a cluster with errors in the event log.

6. Issue the following CLI command to check the status of the software upgrade process:

```
svcinfolsoftwareupgradestatus
```

Note: If a status of `stalled_non_redundant` is displayed, proceeding with the remaining set of node upgrades might result in offline volumes. Contact an IBM service representative to complete the upgrade.

7. Perform the following steps to verify that the software upgrade successfully completed:

- a. Issue the **svctask dumperrlog** CLI command to send the contents of the event log to a text file.

If the software is successfully upgraded, the following output is displayed in the text file:

```
Upgrade completed successfully
```

- b. Issue the **svcinfolnodevdpd** CLI command for each node that is in the cluster. The software version field displays the new software level.

When a new software level is applied, it is automatically installed on all the nodes that are in the cluster.

Note: The software upgrade can take up to 30 minutes per node.

Chapter 4. Audit log commands

An audit log keeps track of action commands that are issued through a Secure Shell (SSH) session or through the management GUI.

The audit log entries provide the following information:

- Identity of the user who issued the action command
- The name of the actionable command
- The timestamp of when the actionable command was issued on the configuration node
- The parameters which were issued with the actionable command

The following commands are not documented in the audit log:

- `svctask dumpconfig`
- `svctask cpdumps`
- `svctask clear.dumps`
- `svctask finderr`
- `svctask dumperrlog`
- `svctask dumpinternallog`
- `svcservicetask dumperrlog`
- `svcservicetask finderr`

The following items are also not documented in the audit log:

- Commands that fail are not logged
- A result code of 0 (success) or 1 (success in progress) is not logged
- Result object ID of node type (for the `addnode` command) is not logged
- Views are not logged

catauditlog

Use the `catauditlog` command to display the in-memory contents of the audit log.

Syntax

```
▶▶— svcinfo — — catauditlog — —————┐—————▶▶  
└— -first — number_of_entries_to_return ┘
```

Parameters

`-first number_of_entries_to_return`

(Optional) Specifies the number of most recent entries to display.

Description

This command lists a specified number of the most recently audited commands.

The in-memory portion of the audit log can hold approximately 1 MB of audit information. Depending on the command text size and the number of parameters, 1 MB records approximately 6000 commands.

Once the in-memory audit log has reached its maximum capacity, the log is written to a local file on the configuration node in the `/dumps/audit` directory. The `svcinfcatauditlog` command only displays the in-memory part of the audit log; the on-disk part of the audit log is in readable text format and does not need any special command to decode it.

The in-memory log entries are reset and cleared automatically, ready to start accumulating new commands. The on-disk portion of the audit log can then be analyzed at a later date.

- | The `lsdumps` command with `-prefix /dumps/audit` can be used to list the files that are on the disk.

The in-memory portion of the audit log can be transferred to an on-disk file using the `svctask dumpauditlog` command. This action clears the in-memory portion of the log.

In the following example, the user has specified that they want to list the five most recent audit log entries.

An invocation example

- | `svcinfcatauditlog -delim : -first 5`

The resulting output

```
| audit_seq_no:timestamp:cluster_user:ssh_ip_address:result:res_obj_id:action_cmd
| 35:091012114520:superuser:9.20.160.249:0::svctask dumpauditlog
| 36:091012115150:superuser:9.20.160.249:0::svctask chquorum -mdisk 45 3
| 37:091012115256:superuser:9.20.160.249:0::svctask chvdisk -name vdisk_master 1
| 38:091012115302:superuser:9.20.160.249:0::svctask chvdisk -name vdisk_aux 2
| 39:091012115328:superuser:9.20.160.249:0::svctask chvdisk -name disk 3
```

dumpauditlog

Use the `dumpauditlog` command to reset or clear the contents of the in-memory audit log. The contents of the audit log are sent to a file in the `/dumps/audit` directory on the current configuration node.

Syntax

▶— `svctask` — — `dumpauditlog` — —▶

Parameters

There are no parameters.

Description

This command dumps the contents of the audit log to a file on the current configuration node. It also clears the contents of the audit log. This command is logged as the first entry in the new audit log.

Audit log dumps are automatically maintained in the `/dumps/audit` directory. The local file system space is used by audit log dumps and is limited to 200 MB on any node in the cluster. The space limit is maintained automatically by deleting the minimum number of old audit log dump files so that the `/dumps/audit` directory space is reduced below 200 MB. This deletion occurs once per day on every node in the cluster. The oldest audit log dump files are considered to be the ones with the lowest audit log sequence number. Also, audit log dump files with a cluster ID number that does not match the current one are considered to be older than files that match the cluster ID, regardless of sequence number.

Other than by running dumps (or copying dump files among nodes), you cannot alter the contents of the audit directory. Each dump file name is generated automatically in the following format:

`auditlog_firstseq_lastseq_timestamp_clusterid`

where

- *firstseq* is the audit log sequence number of the first entry in the log
- *lastseq* is the audit sequence number of the last entry in the log
- *timestamp* is the timestamp of the last entry in the audit log that is being dumped
- *clusterid* is the cluster ID at the time that the dump was created

The audit log dump files names cannot be changed.

The audit log entries in the dump files contain the same information as displayed by the **svcinfo catauditlog** command; however, the **svctask dumpauditlog** command displays the information with one field per line. The **svcinfo lsauditlogdumps** command displays a list of the audit log dumps that are available on the nodes in the cluster.

An invocation example

```
svctask dumpauditlog
```

The resulting output

No feedback

lsauditlogdumps (Deprecated)

- | **Attention:** The **svcinfo lsauditlogdumps** command is deprecated. Use the **svcinfo lsdumps** command to
- | display a list of files in a particular dumps directory.

Chapter 5. User management commands

You can use the command-line interface (CLI) to configure remote authentication service and manage users and user groups on the cluster.

The following user management commands are available:

chauthservice

Configures the remote authentication service of the cluster.

chcurrentuser

Changes the attributes of the current user.

chuser Changes the attributes of an existing user.

chusergrp

Changes the attributes of an existing user group.

lscurrentuser

Displays the name and role of the logged-in user. See the “lscurrentuser” on page 274 command in the Chapter 22, “Information commands,” on page 263 section.

lsuser Lists the users who have been created on the cluster. See the “lsuser” on page 349 command in the Chapter 22, “Information commands,” on page 263 section.

lsusergrp

Lists the user groups that have been created on the cluster. See the “lsusergrp” on page 351 command in the Chapter 22, “Information commands,” on page 263 section.

mkuser

Creates either a local or a remote user to access a SAN Volume Controller cluster.

mkusergrp

Creates a new user group.

rmuser

Removes a user.

rmusergrp

Removes a user group.

chauthservice

The **chauthservice** command can be used to configure the remote authentication service of the cluster.

Syntax

```
▶▶ svctask -- chauthservice -- [-enable {yes|no}] [-url url]
▶ [-username user_name] [-password password] [-sslcert file_name]
▶ [-refresh]
```

Parameters

-enable *yes | no*

(Optional) Enables or disables the SAN Volume Controller system's use of the remote authentication server. When the **enable** parameter is set to **no**, remote authentications are failed by the system, but local authentications continue to operate normally.

-url *url*

(Optional) Specifies the Web address of the remote authentication service. This must be a valid IPv4 or IPv6 network address. You can use the following characters: a - z, A - Z, 0 - 9, _, ~, :, [,], %, or /. The maximum length of the Web address is 100 characters.

-username *user_name*

(Optional) Specifies the HTTP basic authentication user name. The user name cannot start or end with a blank. The user name can consist of a string of 1 - 64 ASCII characters with the exception of the following characters: %:",*".

-password *password*

(Optional) Specifies the HTTP basic authentication user password. The password cannot start or end with a blank. It must consist of a string of 6 - 64 printable ASCII characters. The *password* variable is optional. If you do not provide a password, the system prompts you and does not display the password that you type.

-sslcert *file_name*

(Optional) Specifies the name of the file that contains the SSL certificate, in privacy enhanced mail (PEM) format, for the remote authentication server.

-refresh

(Optional) Causes the SAN Volume Controller to invalidate any remote user authorizations that are cached on the cluster. Use this when you modify user groups on the authentication service and want the change to immediately take effect on the SAN Volume Controller.

Description

This command sets the attributes of the remote authentication service on the cluster. It is not necessary to disable the remote authentication service to change its attributes. To disable the remote authentication service in a controlled manner when it is not available, use the **enable** parameter with the **no** option.

When the authentication service is enabled or the configuration is changed, the cluster does not test whether the remote authentication system is operating correctly. To establish whether the system is operating correctly, issue the command **svcinfo lscurrentuser** for a remotely authenticated user. If the output lists the user's roles obtained from the remote authentication server, remote authentication is operating successfully. If the output is an error message, remote authentication is not working correctly, and the error message describes the problem.

If you are using the **url** parameter, the Web address can have either of the following formats:

- `http://network_address:http remote authentication service port number/path_to_service`
- `https://network_address:https remote authentication service port number/path_to_service`

The network address must be an IPv4 or IPv6 address. Do not use the corresponding host name.

For example, if the system network IPv4 address is 9.71.45.108, you could enter either of the following corresponding addresses:

```
http://9.71.45.108:16310/TokenService/services/Trust
https://9.71.45.108:16311/TokenService/services/Trust
```

An invocation example

To fully configure and enable the authentication service:

```
svctask chauthservice -url https://9.71.45.108:16311/TokenService/services/Trust  
-sslcert /tmp/sslCACert.pem -username admin -password password -enable yes
```

The resulting output

No feedback

An invocation example

To disable remote authentication:

```
svctask chauthservice -enable no
```

The resulting output

No feedback

An invocation example

To switch to an HTTPS connection to the authentication service:

```
svctask chauthservice -url https://9.71.45.108:16311/TokenService/services/Trust  
-sslcert /tmp/ssl_cert.pem
```

The resulting output

No feedback

An invocation example

To refresh the SAN Volume Controller remote authorization cache:

```
svctask chauthservice -refresh
```

The resulting output

No feedback

chcurrentuser

The **chcurrentuser** command changes the attributes of the current user.

Syntax

```
▶▶ svctask — — chcurrentuser — [ -password cleartext_password ] [ -nopassword ]  
▶ [ -keyfile sshkey_filename ] [ -nokey ]
```

Parameters

-password *cleartext_password*

(Optional) Specifies the new password to be associated with the current user. The password cannot start or end with a blank. It must consist of a string of 6 - 64 printable ASCII characters. You can optionally specify the password with the **password** parameter. If you do not specify the password, the system prompts you for it before running the command and does not display the password that you type. Either the **password** parameter or the **nopassword** parameter can be set.

-nopassword
(Optional) Specifies that the user's password is to be deleted.

-keyfile *sshkey_filename*
(Optional) Specifies the name of the file that contains the Secure Shell (SSH) public key. Either the **keyfile** parameter or the **nokey** parameter can be set.

-nokey
(Optional) Specifies that the user's SSH key is to be deleted.

Description

Use the **chcurrent user** command to modify the attributes of the current user.

An invocation example

```
svctask chcurrentuser -password secret -nokey
```

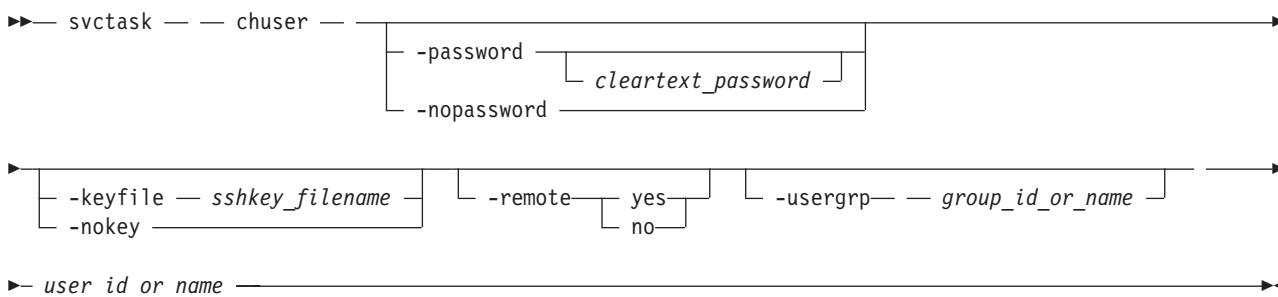
The resulting output

No feedback

chuser

The **chuser** command changes the attributes of an existing user.

Syntax



Parameters

-password *cleartext_password*
(Optional) Specifies the new password to be associated with the user. The password cannot start or end with a blank. It must consist of a string of 6 - 64 printable ASCII characters. You can optionally specify the password with the **password** parameter. If you do not specify the password, the system prompts you for it before running the command and does not display the password that you type. Either the **password** parameter or the **nopassword** parameter can be set.

-nopassword
(Optional) Specifies that the user's password is to be deleted.

-keyfile *sshkey_filename*
(Optional) Specifies the name of the file that contains the Secure Shell (SSH) public key. Either the **keyfile** parameter or the **nokey** parameter can be set.

-nokey
(Optional) Specifies that the user's SSH key is to be deleted.

-remote *yes | no*
(Optional) Specifies whether the user authenticates to the cluster using a remote authentication service. Either the *yes* or *no* option must be set.

-usergrp *group_id_or_name*
(Optional) Specifies the new group for the user.

user_id_or_name
(Required) Specifies the user whose attributes are to be changed.

Description

Use the **chuser** command to modify the attributes of an existing user.

You must have the Security Administrator role to create, delete, or change a user.

Only use the **usergrp** parameter for local users. If you change a user from local to remote, the user's association with any group is removed.

If you change a user from remote to local, a user group must be specified. If you change a user from local to remote, the user must have both a password and an SSH key.

If you use the **keyfile** parameter, the SSH key file should be placed in the /tmp directory before running this command. When you run the command, the SSH key is copied into cluster state and activated for the user, and the input file is deleted.

An invocation example

```
svctask chuser -remote no -usergrp Monitor -nokey jane
```

The resulting output

No feedback

chusergrp

The **chusergrp** command changes the attributes of an existing user group.

Syntax

```
▶▶ svctask — — chusergrp — [ -role — role_name ] [ -remote — yes | no ]  
▶ groupid_or_name —————▶▶
```

Parameters

-role *role_name*
(Optional) Specifies the role to be associated with users that belong to this group. One of the following roles must be selected: Monitor, CopyOperator, Service, Administrator, or SecurityAdmin.

-remote **yes** | **no**
(Optional) Specifies whether this user group should be used to set the role of remote users. Either the yes or no option must be set.

groupid_or_name
(Required) Specifies the user group whose attributes are to be changed.

Description

Use the **chusergrp** command to modify the attributes of an existing user group.

You must have the Security Administrator role to create, delete, or change a user.

The roles of the default groups cannot be changed.

An invocation example

```
svctask chusergrp -role Administrator admin
```

The resulting output

No feedback

mkuser

The **mkuser** command creates either a local or a remote user to access a SAN Volume Controller cluster.

Syntax

```
svctask -- mkuser -- -name -- user_name [ -remote | -usergrp -- group_name_or_id ]
[ -password [ cleartext_password ] | -keyfile -- sshkey_filename ]
```

Parameters

-name *user_name*

(Required) Specifies the unique user name. The user name cannot start or end with a blank. The user name must consist of a string of 1 - 256 ASCII characters, with the exception of the following characters: %:"*'

-remote | **-usergrp** *group_name_or_id*

(Required) Specifies whether the user authenticates to the cluster using a remote authentication service or cluster authentication methods. The **usergrp** parameter specifies the name or ID of the user group with which the local user is to be associated. Either the **remote** parameter or the **usergrp** parameter must be set.

-password *cleartext_password*

(Optional) Specifies the password to be associated with the user. The password cannot start or end with a blank. It must consist of a string of 6 - 64 printable ASCII characters. You can optionally specify the password with the **password** parameter. If you do not specify the password, the system prompts you for it before running the command and does not display the password that you type.

-keyfile *sshkey_filename*

(Optional) Specifies the name of the file that contains the Secure Shell (SSH) public key.

Description

The **mkuser** command creates a new local or remote user to access a cluster. The command returns the ID of the created user.

You must have the Security Administrator role to create, delete, or change a user.

If you create a local user, you must specify the existing user group that the user belongs to. All local users must have a group. The user group defines roles that provide the user with access to specific operations on the cluster. You must also specify either the **keyfile** or **password** parameter, or both.

If you create a remote user, you must specify both the **keyfile** and **password** parameters. Remote users have their groups defined by the remote authentication service.

If you use the **keyfile** parameter, the SSH key file should be placed in the /tmp directory before running this command. When you run the command, the SSH key is copied into cluster state and activated for the user, and the input file is deleted.

An invocation example

```
svctask mkuser -name jane -usergrp Service -password secret
```

The resulting output

```
User, id [1], successfully created
```

mkusergrp

The **mkusergrp** command creates a new user group.

Syntax

```
▶▶ svctask — — mkusergrp — — -name — group_name — -role — role_name — [ -remote ] —▶▶
```

Parameters

-name *group_name*

(Required) Specifies the unique user group name. The group name cannot start or end with a blank. The group name must consist of a string of 1 - 64 ASCII characters, with the exception of the following characters: %:";*' .

-role *role_name*

(Required) Specifies the role to be associated with all users that belong to this user group. One of the following roles must be selected: Monitor, CopyOperator, Service, Administrator, or SecurityAdmin.

-remote

(Optional) Specifies that this user group should be used to set the role of remote users. This is disabled by default.

Description

The **mkusergrp** command creates a new user group to organize users of the SAN Volume Controller cluster by role. Use the **lsusergrp** command to view a list of user groups that have been created on the cluster.

| You must have the security administrator role (SecurityAdmin role name) to create, delete, or change a user group.

Each user group has one role that determines the role of users that belong to that group. Use the **role** parameter to specify one of the following roles for the user group:

Monitor

Users can issue the following commands: `svctask finderr`, `dumperrlog`, `dumpinternallog`, `chcurrentuser`, `ping`, `svconfig backup`, and all of the information commands.

| CopyOperator

Users can issue the following commands: `prestartfcconsistgrp`, `startfcconsistgrp`, `stopfcconsistgrp`, `chfcconsistgrp`, `prestartfcmap`, `startfcmap`, `stopfcmap`, `chfcmap`, `startcrconsistgrp`, `stopcrconsistgrp`, `switchcrconsistgrp`, `chrconsistgrp`, `startcrrelationship`, `stopcrrelationship`, `switchcrrelationship`, `chrrelationship`, and `chpartnership`. In addition, users can issue all of the commands allowed by the Monitor role.

Service

Users can issue the following commands: `applysoftware`, `setlocale`, `addnode`, `rmnode`, `cherrstate`, `writesernum`, `detectmdisk`, `includemdisk`, `clearerrlog`, `cleardumps`, `settimezone`, `stopcluster`, `startstats`, `stopstats`, and `settime`. In addition, users can issue all of the commands allowed by the Monitor role.

Administrator

Users can issue all commands except: `chauthservice`, `mkuser`, `rmuser`, `chuser`, `mkusergrp`, `rmusergrp`, `chusergrp`, and `setpwdreset`

| SecurityAdmin

Users can issue all commands.

The command returns the ID of the created user group.

An invocation example

```
svctask mkusergrp -name support -role Service
```

The resulting output

```
User Group, id [5], successfully created
```

rmuser

The **rmuser** command deletes a user.

Syntax

```
▶— svctask — — rmuser — — user_id_or_name —————▶◀
```

Parameters

user_id_or_name

(Required) Specifies the user to be removed.

Description

Use the **rmuser** command to delete a user.

You must have the Security Administrator role to create, delete, or change a user.

An invocation example

```
svctask rmuser jane
```

The resulting output

```
No feedback
```

rmusergrp

The **rmusergrp** command deletes a user group.

Syntax

```
svctask - - rmusergrp - [-force] group_id_or_name
```

Parameters

-force

(Optional) Specifies that the user group should be deleted even if there are users in the group.

group_id_or_name

(Required) Specifies the ID or name of the user group to be removed.

Description

Use the **rmusergrp** command to delete a user group.

You must have the Security Administrator role to create, delete, or change a user group.

User groups with users cannot normally be deleted. If you use the **force** parameter, the group is deleted and all of the users in that group are assigned to the Monitor group. Default user groups cannot be deleted, even if the **force** parameter is set.

An invocation example

```
svctask rmusergrp support
```

The resulting output

No feedback

Chapter 6. Email and event notification commands

You can use the command-line interface (CLI) to enable your system to send notifications.

The following email and event notification commands are available:

chcluster

Enables the inventory notification function by specifying how often notifications are sent to recipients. See the “chcluster” on page 111 command in the Chapter 7, “Cluster commands,” on page 107 section.

chemail

Sets or modifies contact information for email event notifications.

chemailserver

Modifies the parameters of an existing email server object.

chemailuser

Allows you to modify the email recipient's information.

chsnmpserver

Modifies the parameters of an existing SNMP server.

chsyslogserver

Modifies the parameters of an existing syslog server.

lscluster

Provides information about the addition of email notification recipients to the cluster. See the “lscluster” on page 264 command in the Chapter 22, “Information commands,” on page 263 section.

lsemailserver

Returns a concise list or a detailed view of email servers that are configured on the cluster. See the “lsemailserver” on page 277 command in the Chapter 22, “Information commands,” on page 263 section.

lsemailuser

Generates a report that lists the users (by user ID or type of user [local or support]) who are receiving email event and inventory notifications. See the “lsemailuser” on page 278 command in the Chapter 22, “Information commands,” on page 263 section.

lssnmpserver

Returns a concise list or a detailed view of SNMP servers that are configured on the cluster. See the “lssnmpserver” on page 346 command in the Chapter 22, “Information commands,” on page 263 section.

lssyslogserver

Returns a concise list or a detailed view of syslog servers that are configured on the cluster. See the “lssyslogserver” on page 347 command in the Chapter 22, “Information commands,” on page 263 section.

mkemailserver

Creates an email server object that describes a remote Simple Mail Transfer Protocol (SMTP) email server.

mkemailuser

Allows you to specify the user, the receiving server address, the type of email event notification, and to enable inventory notification in the email event notification facility.

mksnmpserver

Creates an SNMP server to receive notifications.

mksyslogserver

Creates a syslog server to receive notifications.

rmailserver

Deletes the specified email server object.

rmailuser

Removes an existing email recipient from the email event notification function.

rmsnmpserver

Deletes the specified SNMP server.

rmsyslogserver

Deletes the specified syslog server.

sendinventoryemail

Sends an inventory email notification to all inventory email recipients.

startemail

Starts the email notification function.

stopemail

Disables the email notification function.

testemail

Sends a test email to a specified user or to all users using the email notification function.

chemail

The **chemail** command can be used to set or modify contact information for e-mail event notifications. To modify settings, at least one of the parameters must be specified.

Syntax

```

▶▶— svctask — — chemail — [ -reply — reply_email_address ] —————▶
▶ [ -contact — contact_name ] [ -primary — primary_telephone_number ] —————▶
▶ [ -alternate — alternate_telephone_number ] [ -location — location ] —————▶
▶ [ -contact2 — contact_name2 ] [ -primary2 — primary_telephone_number2 ] —————▶
▶ [ -alternate2 — alternate_telephone_number2 ] [ -nocontact2 ] —————▶▶

```

Parameters

-reply *reply_email_address*

(Optional) Specifies the e-mail address to which a reply is sent.

-contact *contact_name*

(Optional) Specifies the name of the person to receive the e-mail.

-primary *primary_telephone_number*

(Optional) Specifies the primary contact telephone number.

- alternate** *alternate_telephone_number*
(Optional) Specifies the alternate contact telephone number that is used when you cannot reach the primary contact on the primary phone.
- location** *location*
(Optional) Specifies the physical location of the system that is reporting the error. The *location* value must not contain punctuation or any other characters that are not alphanumeric or spaces.
- | **-contact2** *contact_name2*
| (Optional) Specifies the name of the second contact person to receive the e-mail.
- | **-primary2** *primary_telephone_number2*
| (Optional) Specifies the primary contact telephone number for the second contact person.
- | **-alternate2** *alternate_telephone_number2*
| (Optional) Specifies the alternate contact telephone number for the second contact person.
- | **-nocontact2**
| (Optional) Removes all the contact details for the second contact person.

Description

This command sets or modifies contact information that is used by the e-mail event notification facility.

Note: If you are starting the e-mail event notification facility, the **reply**, **contact**, **primary**, and **location** parameters are required. If you are modifying contact information used by the e-mail event notification facility, at least one of the parameters must be specified.

An invocation example

```
svctask chemail -primary 0441234567 -location 'room 256 floor 1 IBM'
```

The resulting output

[No feedback]

chemailserver

The **chemailserver** command modifies the parameters of an existing e-mail server object.

Syntax

```

▶▶ svctask — — chemailserver — [ -name — server_name ] [ -ip — ip_address ] —————▶
▶ [ -port — port ] [ email_server_name ] [ email_server_id ] —————▶

```

Parameters

- name** *server_name*
(Optional) Specifies a unique name to assign to the e-mail server object. The name must be a 1-through 15-character string, and cannot start with a hyphen or number. When specifying a server name, *emailserver* is a reserved word.
- ip** *ip_address*
(Optional) Specifies the IP address of the e-mail server object. This must be a valid IPv4 or IPv6 address. IPv6 addresses can be zero compressed.

-port *port*

(Optional) Specifies the port number for the e-mail server. This must be a value of 0 - 65535. The default value is 25.

email_server_name | *email_server_id*

(Required) Specifies the name or ID of the server object to be modified.

Description

Use this command to change the settings of an existing e-mail server object. The e-mail server object describes a remote Simple Mail Transfer Protocol (SMTP) e-mail server.

You must specify either the current name or the ID of the object returned at creation time. Use the **svcinfo lsmailserver** command to obtain this ID.

An invocation example

```
svctask chemailserver -name newserver 0
```

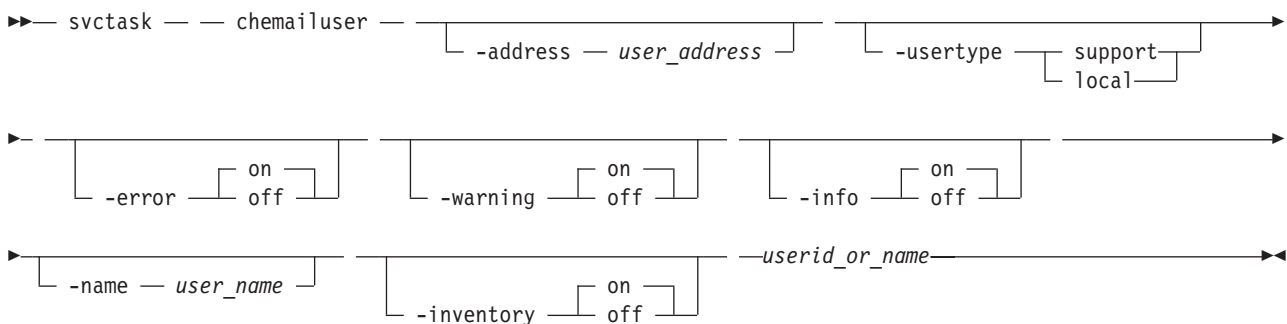
The resulting output

```
none
```

chemailuser

The **chemailuser** command modifies the settings that are defined for an e-mail recipient.

Syntax



Parameters

-address *user_address*

(Optional) Specifies the e-mail address of the person receiving the e-mail or inventory notifications, or both. The *user_address* value must be unique.

-usertype **support** | **local**

(Optional) Specifies the type of user, either local or support, based on the following definitions:

support

Address of the support organization that provides vendor support.

local All other addresses.

-error **on** | **off**

(Optional) Specifies whether the recipient receives error-type event notifications. Set to on, error-type event notifications are sent to the e-mail recipient. Set to off, error-type event notifications are not sent to the recipient.

-warning on | off

(Optional) Specifies whether the recipient receives warning-type event notifications. Set to on, warning-type event notifications are sent to the e-mail recipient. Set to off, warning-type event notifications are not sent to the recipient.

-info on | off

(Optional) Specifies whether the recipient receives informational event notifications. Set to on, informational event notifications are sent to the e-mail recipient. Set to off, informational event notifications are not sent to the recipient.

-name *user_name*

(Optional) Specifies the user name of the new e-mail event notification recipient. The *user_name* value must be unique, must not contain spaces, and must not contain all numbers. The name **emailuser*n***, where *n* is a number, is reserved and cannot be specified as one of your user names.

-inventory on | off

(Optional) Specifies whether this recipient receives inventory e-mail notifications.

userid_or_name

(Required) Specifies the e-mail recipient for whom you are modifying settings.

Description

This command modifies the settings that are established for an e-mail recipient. Standard rules regarding names apply; therefore, it is not possible to change a name to **emailuser*n***, where *n* is a number.

Note: Before the **usertype** parameter can be set to support, the **-warning** and **-info** flags must be set to off.

An invocation example

The following example modifies e-mail settings for e-mail recipient **manager2008**:

```
svctask chemailuser -usertype local manager2008
```

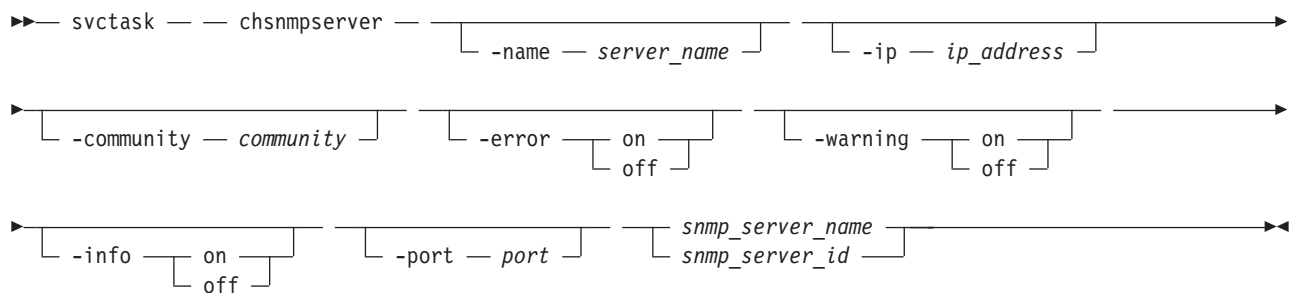
The resulting output

[no feedback]

chsnmpserver

The **chsnmpserver** command modifies the parameters of an existing SNMP server.

Syntax



Parameters

-name *server_name*

(Optional) Specifies a name to assign to the SNMP server. The name must be unique. When specifying a server name, *snmp* is a reserved word.

-ip *ip_address*

(Optional) Specifies an IP address to assign to the SNMP server. This must be a valid IPv4 or IPv6 address.

-community *community*

(Optional) Specifies the community name for the SNMP server.

-error on | off

(Optional) Specifies whether the server receives error notifications. Set to *on*, error notifications are sent to the SNMP server. Set to *off*, error notifications are not sent to the SNMP server.

-warning on | off

(Optional) Specifies whether the server receives warning notifications. Set to *on*, warning notifications are sent to the SNMP server. Set to *off*, warning notifications are not sent to the SNMP server.

-info on | off

(Optional) Specifies whether the server receives information notifications. Set to *on*, information notifications are sent to the SNMP server. Set to *off*, information notifications are not sent to the SNMP server.

-port *port*

(Optional) Specifies the remote port number for the SNMP server. This must be a value of 1 - 65535.

snmp_server_name | *snmp_server_id*

(Required) Specifies the name or ID of the server to be modified.

Description

Use this command to change the settings of an existing SNMP server. You must specify either the current name of the server or the ID returned at creation time. Use the **svcinfolssnmpserver** command to obtain this ID.

An invocation example

```
svctask chsnmpserver -name newserver 0
```

The resulting output


```
none
```

chsyslogserver

The **chsyslogserver** command modifies the parameters of an existing syslog server.

Syntax

```
►► svctask -- chsyslogserver -- [ -name -- server_name ] [ -ip -- ip_address ]
[ -facility -- facility ] [ -error [ on | off ] ] [ -warning [ on | off ] ]
```



Parameters

-name *server_name*

(Optional) Specifies a name to assign to the syslog server. The name must be unique. When specifying a server name, `syslog` is a reserved word.

-ip *ip_address*

(Optional) Specifies an IP address to assign to the syslog server. This must be a valid IPv4 or IPv6 address.

-facility *facility*

(Optional) Specifies a facility number to identify the origin of the message to the receiving server. Servers configured with facility values of 0 - 3 receive syslog messages in concise format. Servers configured with facility values of 4 - 7 receive syslog messages in fully-expanded format.

-error on | off

(Optional) Specifies whether the server receives error notifications. Set to `on`, error notifications are sent to the syslog server. Set to `off`, error notifications are not sent to the syslog server.

-warning on | off

(Optional) Specifies whether the server receives warning notifications. Set to `on`, warning notifications are sent to the syslog server. Set to `off`, warning notifications are not sent to the syslog server.

-info on | off

(Optional) Specifies whether the server receives information notifications. Set to `on`, information notifications are sent to the syslog server. Set to `off`, information notifications are not sent to the syslog server.

syslog_server_name | *syslog_server_id*

(Required) Specifies the name or ID of the server to be modified.

Description

Use this command to change the settings of an existing syslog server. You must specify either the current name of the server or the ID returned at creation time. Use the `svcinfolssyslogserver` command to obtain this ID.

An invocation example

```
svctask chsyslogserver -facility 5 2
```

The resulting output

```
none
```

mkemailserver

The `mkemailserver` command creates an e-mail server object that describes a remote Simple Mail Transfer Protocol (SMTP) e-mail server.

Syntax



┌ -port — port ┐

Parameters

-name *server_name*

(Optional) Specifies a unique name to assign to the e-mail server object. The name must be a 1-through 15-character string, and cannot start with a hyphen or number. If a name is not specified, then a system default of `emailservern` is applied, where *n* is the object ID. When specifying a server name, `emailserver` is a reserved word.

-ip *ip_address*

(Required) Specifies the IP address of a remote e-mail server. This must be a valid IPv4 or IPv6 address. IPv6 addresses can be zero compressed.

-port *port*

(Optional) Specifies the port number for the e-mail server. This must be a value of 1 - 65535. The default value is 25.

Description

This command creates an e-mail server object that represents the SMTP server. The SAN Volume Controller uses the e-mail server to send event notification and inventory e-mails to e-mail users. It can transmit any combination of error, warning, and informational notification types.

The SAN Volume Controller supports up to six e-mail servers to provide redundant access to the external e-mail network. The e-mail servers are used in turn until the e-mail is successfully sent from the SAN Volume Controller. The attempt is successful when the SAN Volume Controller gets a positive acknowledgement from an e-mail server that the e-mail has been received by the server.

An invocation example

```
svctask mkemailserver -ip 2.2.2.2 -port 78
```

The resulting output

```
Emailserver id [2] successfully created
```

mkemailuser

The **mkemailuser** command adds a recipient of e-mail event and inventory notifications to the e-mail event notification facility. You can add up to twelve recipients, one recipient at a time.

Syntax

▶▶ svctask — — mkemailuser — ┌ -name — user_name ┐ — -address — user_address — —————▶

▶ -usertype ┌ support ┐ ┌ -error ┐ ┌ on ┐ ┌ -warning ┐ ┌ on ┐
└ local ┘ └ off ┘ └ off ┘ └ off ┘

▶ ┌ -info ┐ ┌ on ┐ ┌ -inventory ┐ ┌ on ┐
└ off ┘ └ off ┘ └ off ┘ └ off ┘

Parameters

-name *user_name*

(Optional) Specifies the name of the person who is the recipient of e-mail event notifications. The *user_name* value must be unique, must not contain spaces, and must not contain only numbers. If you do not specify a user name, the system automatically assigns a user name in the format of *emailusern*, where *n* is a number beginning with 0 (*emailuser0*, *emailuser1*, and so on).

The name **emailusern**, where *n* is a number, is reserved and cannot be used as one of your user names.

-address *user_address*

(Required) Specifies the e-mail address of the person receiving the e-mail event or inventory notifications, or both. The *user_address* value must be unique.

-usertype **support** | **local**

(Required) Specifies the type of user, either customer or support, based on the following definitions:

support

Address of the support organization that provides vendor support.

local All other addresses.

-error **on** | **off**

(Optional) Specifies whether the recipient receives error-type event notifications. Set to **on**, error-type event notifications are sent to the e-mail recipient. Set to **off**, error-type event notifications are not sent to the recipient. The default value is **on**.

-warning **on** | **off**

(Optional) Specifies whether the recipient receives warning-type event notifications. Set to **on**, warning-type event notifications are sent to the e-mail recipient. Set to **off**, warning-type event notifications are not sent to the recipient. The default value is **on**.

-info **on** | **off**

(Optional) Specifies whether the recipient receives informational event notifications. Set to **on**, informational event notifications are sent to the e-mail recipient. Set to **off**, informational event notifications are not sent to the recipient. The default value is **on**.

-inventory **on** | **off**

(Optional) Specifies whether this recipient receives inventory e-mail notifications. The default value is **off**.

Description

This command adds e-mail recipients to the e-mail event and inventory notification facility. You can add up to twelve recipients, one recipient at a time. When an e-mail user is added, if a user name is not specified, a default name is allocated by the system. This default name has the form of *emailuser1*, *emailuser2*, and so on. E-mail notification starts when you process the **startemail** command.

Note: Before the **usertype** parameter can be set to **support**, the **-warning** and **-info** flags must be set to **off**.

An invocation example

```
svctask mkemailuser -address manager2008@ibm.com -error on -usertype local
```

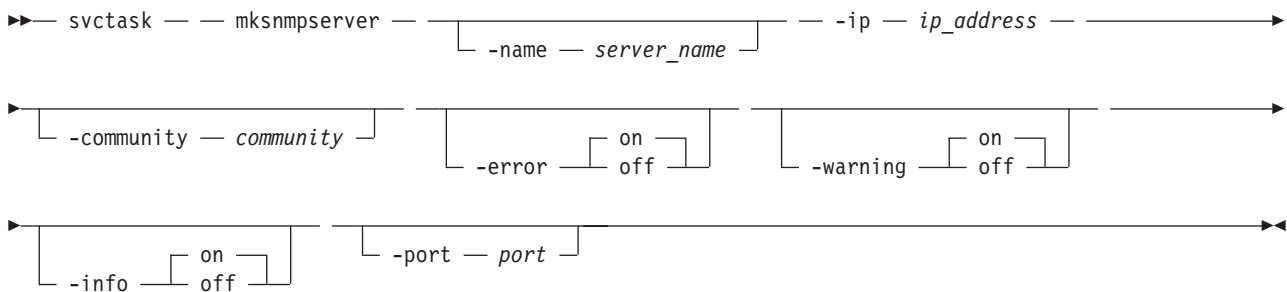
The resulting output

```
email user, id [2], successfully created
```

mkxsnmpserver

The **mkxsnmpserver** command creates an SNMP server to receive notifications.

Syntax



Parameters

-name *server_name*

(Optional) Specifies a unique name to assign to the SNMP server. If a name is not specified, then a system default of `snmpn` is applied, where *n* is the ID of the server. When specifying a server name, `snmp` is a reserved word.

-ip *ip_address*

(Required) Specifies the IP address of the SNMP server. This must be a valid IPv4 or IPv6 address.

-community *community*

(Optional) Specifies the community name for the SNMP server. If you do not specify a community name, then the default name of `public` is used.

-error on | off

(Optional) Specifies whether the server receives error notifications. Set to `on`, error notifications are sent to the SNMP server. Set to `off`, error notifications are not sent to the SNMP server. The default value is `on`.

-warning on | off

(Optional) Specifies whether the server receives warning notifications. Set to `on`, warning notifications are sent to the SNMP server. Set to `off`, warning notifications are not sent to the SNMP server. The default value is `on`.

-info on | off

(Optional) Specifies whether the server receives information notifications. Set to `on`, information notifications are sent to the SNMP server. Set to `off`, information notifications are not sent to the SNMP server. The default value is `on`.

-port *port*

(Optional) Specifies the remote port number for the SNMP server. This must be a value of 1 - 65535. The default value is 162.

Description

This command creates an SNMP server to receive notifications.

SAN Volume Controller supports a maximum of 6 SNMP servers.

An invocation example

```
svctask mksnmpserver -ip 2.2.2.2 -port 78
```

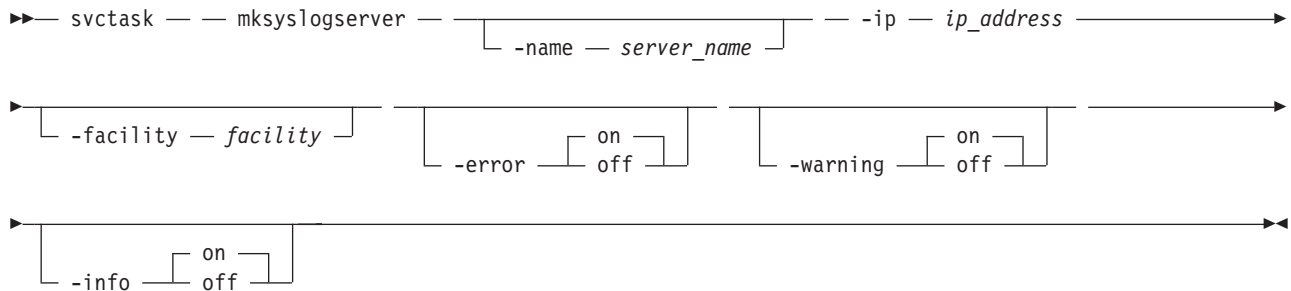
The resulting output

```
SNMP Server id [2] successfully created
```

mksyslogserver

The `mksyslogserver` command creates a syslog server to receive notifications.

Syntax



Parameters

-name *server_name*

(Optional) Specifies a unique name to assign to the syslog server. If a name is not specified, then a system default of `syslog n` is applied, where n is the ID of the server. When specifying a server name, `syslog` is a reserved word.

-ip *ip_address*

(Required) Specifies the IP address of the syslog server. This must be a valid IPv4 or IPv6 address.

-facility *facility*

(Optional) Specifies the facility number used in syslog messages. This number identifies the origin of the message to the receiving server. Servers configured with facility values of 0 - 3 receive syslog messages in concise format. Servers configured with facility values of 4 - 7 receive syslog messages in fully-expanded format. The default value is 0.

-error on | off

(Optional) Specifies whether the server receives error notifications. Set to `on`, error notifications are sent to the syslog server. Set to `off`, error notifications are not sent to the syslog server. The default value is `on`.

-warning on | off

(Optional) Specifies whether the server receives warning notifications. Set to `on`, warning notifications are sent to the syslog server. Set to `off`, warning notifications are not sent to the syslog server. The default value is `on`.

-info on | off

(Optional) Specifies whether the server receives information notifications. Set to `on`, information notifications are sent to the syslog server. Set to `off`, information notifications are not sent to the syslog server. The default value is `on`.

Description

This command creates a syslog server to receive notifications. The syslog protocol is a client-server standard for forwarding log messages from a sender to a receiver on an IP network. Syslog can be used to integrate log messages from different types of systems into a central repository.

SAN Volume Controller supports a maximum of 6 syslog servers.

An invocation example

```
svctask mksyslogserver -ip 1.2.3.4
```

The resulting output

Syslog Server id [2] successfully created

rmemailserver

The **rmemailserver** command deletes the specified e-mail server object.

Syntax

```
svctask -- rmemailserver -- email_server_name | email_server_id
```

Parameters

email_server_name | *email_server_id*

(Required) Specifies the name or ID of the e-mail server object to be deleted.

Description

Use this command to delete an existing e-mail server object that describes a remote Simple Mail Transfer Protocol (SMTP) e-mail server. You must specify either the current name or the ID of the object returned at creation time. Use the **svctask lsemailserver** command to obtain this ID.

Note: E-mail service stops when the last e-mail server is removed. Use the **svctask startemail** command to reactivate the e-mail and inventory notification function after at least one e-mail server has been configured.

An invocation example

```
svctask rmemailserver email4
```

The resulting output

none

rmemailuser

The **rmemailuser** command allows you to remove a previously defined e-mail recipient from your system.

Syntax

```
svctask -- rmemailuser -- userid_or_name
```

Parameters

userid_or_name

(Required) Specifies the user ID or user name of the e-mail recipient to remove.

Description

This command removes an existing e-mail recipient from the system.

An invocation example

The following example removes e-mail recipient **manager2008**:


```
svctask rmemailuser manager2008
```

The resulting output

[No feedback]

An invocation example

The following example removes e-mail recipient 2:

```
svctask rmemailuser 2
```

The resulting output

[No feedback]

rmsnmpserver

The **rmsnmpserver** command deletes the specified SNMP server.

Syntax

```
▶▶ svctask — — rmsnmpserver — [ snmp_server_name ] —————▶▶  
                               [ snmp_server_id ]
```

Parameters

snmp_server_name | *snmp_server_id*

(Required) Specifies the name or ID of the SNMP server to be deleted.

Description

Use this command to delete an existing SNMP server. You must specify either the current name of the server or the ID returned at creation time. Use the **svcinfo lssnmpserver** command to obtain this ID.

An invocation example

```
svctask rmsnmpserver snmp4
```

The resulting output

none

rmsyslogserver

The **rmsyslogserver** command deletes the specified syslog server.

Syntax

```
▶▶ svctask — — rmsyslogserver — [ syslog_server_name ] —————▶▶  
                               [ syslog_server_id ]
```

Parameters

syslog_server_name | *syslog_server_id*

(Required) Specifies the name or ID of the syslog server to be deleted.

Description

Use this command to delete an existing syslog server. You must specify either the current name of the server or the ID returned at creation time. Use the `svcinfolssyslogserver` command to obtain this ID.

An invocation example

```
svctask rmysyslogserver 2
```

The resulting output

```
none
```

sendinventoryemail

The `sendinventoryemail` command sends an inventory e-mail notification to all e-mail recipients who are enabled to receive inventory e-mail notifications. There are no parameters for this command.

Syntax

```
svctask — — sendinventoryemail
```

Parameters

There are no parameters for this command.

Description

This command sends an inventory e-mail notification to all e-mail recipients who are enabled to receive inventory e-mail notifications. This command fails if the `startemail` command has not been processed and at least one e-mail recipient using the e-mail event and inventory notification facility has not been set up to receive inventory e-mail notifications. This command also fails if the e-mail infrastructure has not been set up.

An invocation example

In the following example, you send an inventory e-mail notification to all e-mail recipients who are enabled to receive them:

```
svctask sendinventoryemail
```

The resulting output

```
[No feedback]
```

startemail

The `startemail` command activates the e-mail and inventory notification function. There are no parameters for this command.

Syntax

```
svctask — — startemail
```

Parameters

There are no parameters for this command.

Description

This command enables the e-mail event notification service. No e-mails are sent to users until the **startemail** command has been run and at least one user has been defined to the system.

An invocation example

In the following example, you are starting the e-mail error notification service.

```
svctask startemail
```

The resulting output

```
[No feedback]
```

stopemail

The **stopemail** command stops the e-mail and inventory notification function. There are no parameters for this command.

Syntax

```
▶▶ svctask — — stopemail —————▶▶
```

Parameters

There are no parameters for this command.

Description

This command stops the e-mail error notification function. No e-mails are sent to users until the **startemail** command is reissued.

An invocation example

In the following example, you have stopped the e-mail and inventory notification function:

```
svctask stopemail
```

The resulting output

```
[No feedback]
```

testemail

The **testemail** command allows you to send an e-mail notification to one user of the e-mail notification function or to all users of the e-mail notification function to ensure that the function is operating correctly.

Syntax

```
▶▶ svctask — — testemail —————▶▶  
                        └──userid_or_name──┘  
                        └──-all──┘
```

Parameters

userid_or_name

(Required if you do not specify **-all**) Specifies the user ID or user name of the e-mail recipient that you want to send a test e-mail to. You cannot use this parameter with the **-all** parameter. The *userid_or_name* value must not contain spaces.

-all

(Required if you do not specify *userid_or_name*) Sends a test e-mail to all e-mail users configured to receive notification of events of any notification type. No attempt is made to send the test e-mail to an e-mail user who does not have any notification setting set to *on*.

Description

This command sends test e-mails to the specified e-mail users. The e-mail recipient expects to receive the test e-mail within a specified service time. If the e-mail is not received within the expected time period, the recipient must contact the administrator to ensure that the e-mail settings for the user are correct. If there is still a problem, you must contact the IBM Support Center.

The e-mail recipient uses the test e-mail to check that the SMTP name, the IP address, the SMTP port, and the user address are valid.

An invocation example

The following example sends a test e-mail to the user ID **manager2008**:

```
svctask testemail manager2008
```

The resulting output

```
[No feedback]
```

Chapter 7. Cluster commands

Cluster commands are used to monitor and modify clusters.

A cluster consists of up to four pairs of nodes that provide a single configuration and service interface. There are a number of cluster commands available for various tasks.

addnode

Adds a new (candidate) node to an existing cluster.

cfgportip

Assigns an IP address to each node ethernet port for iSCSI I/O.

chcluster

Modifies the attributes of an existing cluster.

chclusterip

Modifies the IP configuration parameters for the cluster.

chiogrp

Modifies the name of an I/O group or the amount of memory that is available for Copy Services or VDisk mirroring.

chnode

Changes the name assigned to a node.

cleardumps

Cleans the various dump directories on a specified node.

cpdumps

Copies dump files from a nonconfiguration node onto the configuration node.

detectmdisk

Rescans the fibre-channel network for new managed disks.

ping Used to diagnose IP configuration problems by checking whether the specified IP address is accessible from the configuration node.

rmnode

Deletes a node from the cluster.

rmportip

Removes an iSCSI IP address from a node ethernet port.

setclustertime

Sets the time for the cluster.

setpwdreset

Changes the status of the password-reset feature for the display panel.

settimezone

Sets the time zone for the cluster.

startstats

Starts the collection of per-node statistics for virtual disks (VDisks), managed disks (MDisks), and nodes.

stopcluster

Shuts down a single node or the entire cluster.

addnode

You can use the **addnode** command to add a new (candidate) node to an existing cluster. You can enter this command any time after a cluster has been created. If you are adding a new node to a cluster, you must ensure that the model type of the new node is supported by the SAN Volume Controller software version of the cluster. If the model type is not supported by the cluster software, you must upgrade the cluster to a software version that supports the model type of the new node.

Syntax

```
svctask -- addnode -- [-panelname -- panel_name] [-wwnodename -- wwnn_arg] [-name -- new_name_arg]
-- [-iogrp -- [iogroup_name | iogroup_id]]
```

Parameters

-panelname *panel_name*

(Required if you do not specify the **-wwnodename** parameter) Specifies the node that you want to add to a cluster by the name that is displayed on the display panel. You cannot use this parameter with the **-wwnodename** parameter.

-wwnodename *wwnn_arg*

(Required if you do not specify the **-panelname** parameter) Specifies the node that you want to add to the cluster by the worldwide node name (WWNN). You cannot use this parameter with the **-panelname** parameter.

-name *new_name_arg*

(Optional) Specifies a name for the node that you want to add to the cluster. You can use this name in subsequent commands to refer to the node, instead of using the node ID.

Note: Node names supplied with the **-name** parameter on the **addnode** and **chnode** commands must not already be in use as node names or as node failover_names.

If you assign a name, this name is displayed as the node name from then on. If you do not assign a name, a default name is used. The default name that is used depends on whether the node is replacing one that has previously been deleted. When a node is deleted, its name is retained in the I/O group as the failover name of its partner node. If no nodes remain in an I/O group, no failover names are retained. Only one failover name can be stored for each node. If you add a node into an I/O group that has a retained failover name and do not specify a node name, the retained failover name is assigned to this node. If you do not specify a name and there is no retained failover name, the name assigned has the format nodeX.

Important: The iSCSI Qualified Name (IQN) for each node is generated using the cluster and node names. If you are using the iSCSI protocol and the target name for this node is already active on its partner node, and iSCSI hosts are attached to it. Adding the node with a different name changes the IQN of this node in the cluster and might require reconfiguration of all iSCSI-attached hosts.

-iogrp *iogroup_name* | *iogroup_id*

(Required) Specifies the I/O group to which you want to add this node.

Description

This command adds a new node to the cluster. You can obtain a list of candidate nodes (those that are not already assigned to a cluster) by typing `svcinfolnodecandidate`.

Before you add a node to the cluster, you must check to see if any of the following conditions are true. If the following conditions exist, failure to follow the procedures that are documented here might result in the corruption of all data that is managed by the cluster.

- Is the new node being used to replace a failed node in the cluster?
- Does the node being added to the cluster use physical node hardware that has been used as a node in another cluster, and are both clusters recognized by the same hosts?

If any of the previous conditions are true, you must take the following actions:

1. Add the node to the same I/O group that it was previously in. You can use the command-line interface command **svcinfo lsnode** or the management GUI to determine the WWNN of the cluster nodes.
2. Shut down all of the hosts that use the cluster, before you add the node back into the cluster.
3. Add the node back to the cluster before the hosts are restarted. If the I/O group information is unavailable or it is inconvenient to shut down and restart all of the hosts that use the cluster, you can do the following:
 - a. On all of the hosts that are connected to the cluster, unconfigure the fibre-channel adapter device driver, the disk device driver, and the multipathing driver before you add the node to the cluster.
 - b. Add the node to the cluster and then reconfigure the fibre-channel adapter device driver, the disk device driver, and multipathing driver.

If you are adding a new node to a cluster, take the following actions:

1. Ensure that the model type of the new node is supported by the SAN Volume Controller software version of the cluster. If the model type is not supported by the cluster software, you must upgrade the cluster to a software version that supports the model type of the new node.
2. Record the node serial number, the WWNN, all WWPNs, and the I/O group to which the node has been added. You might need to use this information later. Having it available can prevent possible data corruption if the node must be removed from and re-added to the cluster.

Other considerations when you add a node to a cluster:

When you add a node to the cluster using the **svctask addnode** command or the cluster GUI, you must confirm whether the node has previously been a member of the cluster. If it has, follow one of these two procedures:

- Add the node to the same I/O group that it was previously in. You can determine the WWNN of the nodes in the cluster using the **svcinfo lsnode** command.
- If you cannot determine the WWNN of the nodes in the cluster, call the support team to add the node back into the cluster without corrupting the data.

When a node is added to a cluster, it displays a state of adding. It can take as long as 30 minutes for the node to be added to the cluster, particularly if the software version of the node has changed.

Attention: If the node remains in the adding state for more than 30 minutes, contact your support representative to assist you in resolving this issue.

| When a node is deleted, its name is retained in an I/O group as the failover name of its partner node. If
| no nodes remain in an I/O group, no failover names are retained. The **addnode** command fails if you
| specify a name that is either an existing node name or a retained failover name. Specify a different name
| for the node being added.

An invocation example

```
svctask addnode -wwnodename 5005076801e08b -iogrp io_grp0
```

The resulting output

Node, id [6], successfully added

cfgportip

The **cfgportip** command assigns an IP address to each node ethernet port for iSCSI I/O.

Syntax

```
svctask -- cfgportip --node [ node_name | node_id ] [ -ip -- ipv4addr | -ip_6 -- ipv6addr ]
[ -gw -- ipv4gw | -gw_6 -- ipv6gw ] [ -mask -- subnet_mask | -prefix_6 -- prefix ] [ -failover ] --port_id
```

Parameters

-node *node_name* | *node_id*

(Required) Specifies which node has the ethernet port that the IP address is being assigned to.

-ip *ipv4addr*

(Required if you do not use **ip_6**) Sets the IPv4 address for the ethernet port. You cannot use this parameter with the **ip_6** parameter.

-ip_6 *ipv6addr*

(Required if you do not use **ip**) Sets the IPv6 address for the ethernet port. You cannot use this parameter with the **ip** parameter.

-gw *ipv4addr*

(Required if you do not use **gw_6**) Sets the IPv4 gateway IP address. You cannot use this parameter with the **gw_6** parameter.

-gw_6 *ipv6gw*

(Required if you do not use **gw**) Sets the IPv6 default gateway address for the port. You cannot use this parameter with the **gw** parameter.

-mask *subnet_mask*

(Required if you do not use **prefix_6**) Sets the IPv4 subnet mask. You cannot use this parameter with the **prefix_6** parameter.

-prefix_6 *prefix*

(Required if you do not use **mask**) Sets the IPv6 prefix. You cannot use this parameter with the **mask** parameter.

-failover

(Optional) Specifies that the IP address belongs to the partner node in the I/O group. If the partner node is offline, the address is configured and presented by this node. When another node comes online in the I/O group, the failover address is presented by that node.

port_id

(Required) Specifies which port (1 or 2) to apply changes to.

Description

This command assigns either an IPv4 or IPv6 address to a specified ethernet port of a node. The IP address is used for iSCSI I/O. Use the **svctask chclusterip** command to assign cluster IP addresses.

For an IPv4 address, the **ip**, **mask**, and **gw** parameters are required. All of the IPv4 IP parameters must be specified to assign an IPv4 address to an ethernet port.

For an IPv6 address, the **ip_6**, **prefix_6**, and **gw_6** parameters are required. All of the IPv6 IP parameters must be specified to assign an IPv6 address to an ethernet port.

Use the **svcinfo lsportip** command with the optional **ethernet_port_id** parameter to list the port IP addresses for the specified port.

An invocation example for IPv4

```
svctask cfgportip -node 1 -ip 9.8.7.1 -gw 9.0.0.1 -mask 255.255.255.0 1
```

The resulting output

No feedback

An invocation example for IPv6

```
svctask cfgportip -node 1 -ip_6 3:3:0:4::0 -gw_6 ffe8::0 -prefix_6 64 2
```

The resulting output

No feedback

chcluster

The **chcluster** command modifies the attributes of an existing cluster. You can enter this command any time after a cluster has been created. All the parameters that are associated with this command are optional. However, you must specify one or more parameters with this command.

Syntax

```
▶▶ svctask — chcluster — [ -name — cluster_name ] [ -speed — fabric_speed ]
▶ [ -alias — id_alias ] [ -invemailinterval — interval ]
▶ [ -gmlinktolerance — link_tolerance ] [ -gmmaxhostdelay — max_host_delay ]
▶ [ -gminterdelaysimulation — inter_cluster_delay_simulation ]
▶ [ -gmintradelaysimulation — intra_cluster_delay_simulation ]
▶ [ -ntpip — ipv4_ntp_ip_address ] [ -ntpip_6 — ipv6_ntp_ip_address ]
▶ [ -isnsip — sns_server_address ] [ -isnsip_6 — ipv6_sns_server_address ]
▶ [ -relationshipbandwidthlimit — bandwidth_in_mbps ] [ -iscsiauthmethod — none | chap ]
▶ [ -chapsecret — chap_secret ] [ -nochapsecret ]
```

Parameters

-name *cluster_name*

(Optional) Specifies a new name for the cluster.

Important: The iSCSI Qualified Name (IQN) for each node is generated using the cluster and node names. If you are using the iSCSI protocol, changing either name also changes the IQN of all of the nodes in the cluster and might require reconfiguration of all iSCSI-attached hosts.

-speed *fabric_speed*

(Optional) Specifies the speed of the fabric to which this cluster is attached. Valid values are 1 or 2 (GB).

Attention: Changing the speed on a running cluster breaks I/O service to the attached hosts. Before changing the fabric speed, stop I/O from active hosts and force these hosts to flush any cached data by unmounting volumes (for UNIX host types) or by removing drive letters (for Windows host types). Some hosts might need to be rebooted to detect the new fabric speed.

-alias *id_alias*

(Optional) Specifies an alternate name that does not change the basic ID for the cluster, but does influence the VDisk_UID of every **vdiskhostmap**, both existing and new. These objects appear to have been created for a cluster whose ID matches the alias. Therefore, changing the cluster alias causes loss of host VDisk access, until each host rescans for VDIs presented by the cluster.

-invemailinterval *interval*

(Optional) Specifies the interval at which inventory e-mails are sent to the designated e-mail recipients. The interval range is 0 to 15. The interval is measured in days. Setting the value to 0 turns the inventory e-mail notification function off.

-gmlinktolerance *link_tolerance*

(Optional) Specifies the length of time, in seconds, for which an inadequate intercluster link is tolerated for a Global Mirror operation. The parameter accepts values from 10 to 400 seconds in steps of 10 seconds. The default is 300 seconds. You can disable the link tolerance by entering a value of zero (0) for this parameter.

-gmmaxhostdelay *max_host_delay*

(Optional) Specifies the maximum time delay, in milliseconds, above which the Global Mirror link tolerance timer starts counting down. This threshold value determines the additional impact that Global Mirror operations can add to the response times of the Global Mirror source VDIs (volumes). You can use this parameter to increase the threshold from the default value of 5 milliseconds.

-gminterdelaysimulation *inter_cluster_delay_simulation*

(Optional) Specifies the intercluster delay simulation, which simulates the Global Mirror round trip delay between two clusters, in milliseconds. The default is 0; the valid range is 0 to 100 milliseconds.

-gmintradelaysimulation *intra_cluster_delay_simulation*

(Optional) Specifies the intracluster delay simulation, which simulates the Global Mirror round trip delay in milliseconds. The default is 0; the valid range is 0 to 100 milliseconds.

-ntpip *ipv4_ntp_ip_address*

(Optional) Specifies the IPv4 address for the Network Time Protocol (NTP) server. Configuring an NTP server address causes the cluster to immediately start using that NTP server as its time source. To stop using the NTP server as a time source, invoke the **-ntpip** parameter with a zero address, as follows:

```
svctask chcluster -ntpip 0.0.0.0
```

-ntpip_6 *ipv6_ntp_ip_address*

(Optional) Specifies the IPv6 address for the NTP server. Configuring an NTP server address causes the cluster to immediately start using that NTP server as its time source. To stop using the NTP server as a time source, invoke the **-ntpip_6** parameter with a zero address, as follows:

```
svctask chcluster -ntpip_6 0::0
```

-isnsip *sns_server_address*

(Optional) Specifies the IPv4 address for the iSCSI storage name service (SNS). To stop using the configured IPv4 iSCSI SNS server, invoke the **-isnsip** parameter with a zero address, as follows:

```
svctask chcluster -isnsip 0.0.0.0
```

-isnsip_6 *ipv6_sns_server_address*

(Optional) Specifies the IPv6 address for the iSCSI SNS. To stop using the configured IPv6 iSCSI SNS server, invoke the **-isnsip_6** parameter with a zero address, as follows:

```
svctask chcluster -isnsip_6 0::0
```

-relationshipbandwidthlimit *bandwidth_in_mbps*

(Optional) Specifies the new background copy bandwidth in megabytes per second (MBps), from 1 - 1000. The default is 25 MBps. This parameter operates cluster-wide and defines the maximum background copy bandwidth that any relationship can adopt. The existing background copy bandwidth settings defined on a partnership continue to operate, with the lower of the partnership and VDisk rates attempted.

Note: Do not set this value higher than the default without establishing that the higher bandwidth can be sustained.

-iscsiauthmethod *none | chap*

(Optional) Sets the authentication method for the iSCSI communications of the cluster. The **iscsiauthmethod** value can be none or chap.

-chapsecret *chap_secret*

(Optional) Sets the Challenge Handshake Authentication Protocol (CHAP) secret to be used to authenticate the cluster via iSCSI. This parameter is required if the **iscsiauthmethod chap** parameter is specified. The specified CHAP secret cannot begin or end with a space.

-nochapsecret

(Optional) Clears any previously set CHAP secret for iSCSI authentication. This parameter is not allowed if the **chapsecret** parameter is specified.

Description

This command modifies specific features of a cluster. Multiple features can be changed by issuing a single command.

Using the **-ntpip** or **-ntpip_6** parameter allows the cluster to use an NTP server as an outside time source. The cluster adjusts the system clock of the configuration node according to time values from the NTP server. The clocks of the other nodes are updated from the configuration node's clock. In the NTP mode, the **svctask setclustertime** command is disabled.

All command parameters are optional; however, you must specify at least one parameter.

Use the **svctask chclusterip** command to modify the cluster IP address and service IP address.

An invocation example

```
svctask chcluster -ntpip 9.20.165.16
```

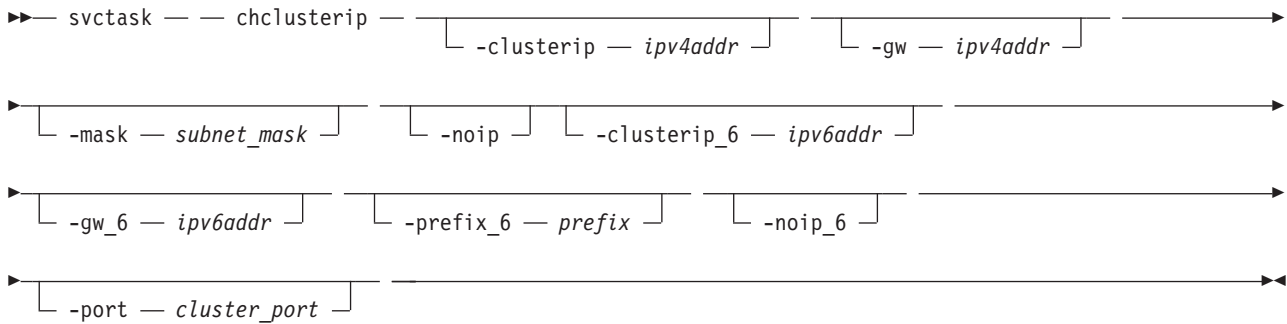
The resulting output

No feedback

chclusterip

The **chclusterip** command modifies the IP configuration parameters for the cluster.

Syntax



Parameters

-clusterip *ipv4addr*

(Optional) Changes the IPv4 cluster IP address. When you specify a new IP address for a cluster, the existing communication with the cluster is broken.

-gw *ipv4addr*

(Optional) Changes the IPv4 default gateway IP address of the cluster.

-mask *subnet_mask*

(Optional) Changes the IPv4 subnet mask of the cluster.

-noip

(Optional) Unconfigures the IPv4 stack on the specified port, or both ports if none is specified.

| **Note:** This parameter does not affect node service address configurations.

-clusterip_6 *ipv6addr*

(Optional) Sets the IPv6 cluster address for the port.

-gw_6 *ipv6addr*

(Optional) Sets the IPv6 default gateway address for the port.

-prefix_6 *prefix*

(Optional) Sets the IPv6 prefix.

-noip_6

(Optional) Unconfigures the IPv6 stack on the specified port, or both ports if none is specified.

| **Note:** This parameter does not affect node service address configurations.

-port *cluster_port*

(Optional) Specifies which port (1 or 2) to apply changes to. This parameter is required unless the **noip** or **noip_6** parameter is used.

Description

This command modifies IP configuration parameters for the cluster. The first time you configure a second port, all IP information is required. Port 1 on the cluster must always have one stack fully configured.

| There are two active cluster ports on the configuration node. There are also two active service ports on any node in which you are performing a service action.

If the cluster IP address is changed, the open command-line shell closes during the processing of the command. You must reconnect to the new IP address if connected through that port.

The **noip** and **noip_6** parameters can be specified together only if the **port** is also specified. The **noip** and **noip_6** parameters cannot be specified with any parameters other than **port**.

- | **Note:** The **noip** and **noip_6** parameters do not affect node service address configurations. Port 1 must have an IPv4 or IPv6 cluster address. The configuration of port 2 is optional.

Service IP addresses for all ports and stacks are initialized to DHCP. A service IP address is always configured.

Modifying an IP address: List the IP address of the cluster by issuing the **svcinfo lscluster** command. Modify the IP address by issuing the **svctask chclusterip** command. You can either specify a static IP address or have the system assign a dynamic IP address.

Table 8 provides IP address formats that are supported.

Table 8. *ip_address_list* formats

IP type	<i>ip_address_list</i> format
IPv4 (no port set, cluster uses default)	1.2.3.4
IPv4 with specific port	1.2.3.4:22
Full IPv6, default port	1234:1234:abcd:0123:0000:0000:7689:6576
Full IPv6, default port, leading zeros suppressed	1234:1234:abcd:123:0:0:7689:6576
Full IPv6 with port	[1234:1234:abcd:0123:0000:0000:7689:6576]:23
Zero-compressed IPv6, default port	1234:1234:abcd:123::7689:6576
Zero-compressed IPv6 with port	[1234:1234:abcd:123::7689:6576]:23

An invocation example

```
svctask chclusterip -clusterip 9.20.136.5 -gw 9.20.136.1 -mask 255.255.255.0 -port 1
```


The resulting output

No feedback

chiogrp

The **chiogrp** command modifies the name of an I/O group, or the amount of memory that is available for Copy Services or VDisk mirroring operations.

Syntax



Parameters

-name *new_name*

(Optional) Specifies the name to assign to the I/O group. The **-name** parameter cannot be specified with the **-feature**, **-size**, or **-kb** parameters.

| **-feature** **flash** | **remote** | **mirror** | **raid**

(Optional) Specifies the feature to modify the amount of memory for: Copy Services or VDisk mirroring. You must specify this parameter with the **-size** parameter. You cannot specify this parameter with the **-name** parameter.

Note: Specifying **remote** changes the amount of memory that is available for Metro Mirror or Global Mirror processing. Any VDisk that is in a Metro Mirror or Global Mirror relationship uses memory in its I/O group, including master and auxiliary VDIs, and VDIs that are in inter-cluster or intra-cluster relationships.

-size *memory_size*

(Optional) Specifies the amount of memory that is available for the specified Copy Services or VDisk mirroring function. Valid input is **0** or any integer. The default unit of measurement for this parameter is megabytes (MB); you can use the kilobytes **-kb** parameter to override the default. You must specify this parameter with the **-feature** parameter. You cannot specify this parameter with the **-name** parameter.

-kb

(Optional) Changes the units for the **-size** parameter from megabytes (MB) to kilobytes (KB). If you specify this parameter, the **-size** *memory_size* value must be any number divisible by 4. You must specify this parameter with the **-feature** and **-size** parameters. You cannot specify this parameter with the **-name** parameter.

io_group_id | *io_group_name*

(Required) Specifies the I/O group to modify. You can modify an I/O group by using the **-name** or the **-feature** parameter.

| **-maintenance** *yes* | *no*

(Optional) Specifies whether the I/O group should be in maintenance mode. The I/O group should be placed in maintenance mode while carrying out service procedures on storage enclosures. Once you enter maintenance mode, it continues until either:

- It is explicitly cleared, OR
- 30 minutes elapse

Description

The **chiogrp** command modifies the name of an I/O group or the amount of memory that is available for Copy Services or VDisk mirroring. You can assign a name to an I/O group or change the name of a specified I/O group. You can change the amount of memory that is available for Copy Services or VDisk mirroring operations by specifying the **-feature** **flash** | **remote** | **mirror** parameter, and a memory size. For VDisk mirroring and Copy Services (FlashCopy, Metro Mirror, and Global Mirror), memory is traded against memory that is available to the cache. The amount of memory can be decreased or increased. Consider the following memory sizes when you use this command:

- The default memory size for FlashCopy is 20 MB.
- The default memory size for Metro Mirror and Global Mirror is 20 MB.
- | • The default memory size for mirrored VDIs is 20 MB.
- The maximum memory size that can be specified for FlashCopy is 512 MB.
- The maximum memory size that can be specified for Metro Mirror and Global Mirror is 512 MB.
- The maximum memory size that can be specified for mirrored VDIs is 512 MB.
- | • The maximum combined memory size across all features is 552 MB.

Table 9 demonstrates the amount of memory required for VDisk mirroring and Copy Services. Each 1 MB of memory provides the following VDisk capacities and grain sizes:

Table 9. Memory required for VDisk Mirroring and Copy Services

Feature	Grain size	1 MB of memory provides the following VDisk capacity for the specified I/O group
Metro Mirror and Global Mirror	256 KB	2 TB of total Metro Mirror and Global Mirror VDisk capacity
FlashCopy	256 KB	2 TB of total FlashCopy source VDisk capacity
FlashCopy	64 KB	512 GB of total FlashCopy source VDisk capacity
Incremental FlashCopy	256 KB	1 TB of total Incremental FlashCopy source VDisk capacity
Incremental FlashCopy	64 KB	256 GB of total Incremental FlashCopy source VDisk capacity
VDisk mirroring	256 KB	2 TB of mirrored VDIs

Table 10 provides an example of RAID level comparisons with their bitmap memory cost, where *MS* is the size of the member drives and *MC* is the number of member drives.

Table 10. RAID level comparisons

Level	Member count	Approximate capacity	Redundancy	Approximate bitmap memory cost
RAID-0	1-8	$MC * MS$	None	(1 MB per 2 TB of MS) * MC
RAID-1	2	MS	1	(1 MB per 2 TB of MS) * (MC/2)
RAID-5	3-16	$(MC-1) * MS$	1	1 MB per 2 TB of MS with a strip size of 256 KB; double with strip size of 128 KB.
RAID-6	5-16	less than $(MC-2 * MS)$	2	
RAID-10	2-16 (evens)	$MC/2 * MS$	1	(1 MB per 2 TB of MS) * (MC/2)

Note: There is a margin of error on the approximate bitmap memory cost of approximately 15%. For example, the cost for a 256 KB RAID-5 is ~1.15 MB for the first 2 TB of MS.

For multiple FlashCopy targets, you must consider the number of mappings. For example, for a mapping with a 256 KB grain size, 8 KB of memory allows one mapping between a 16 GB source VDisk and a 16 GB target VDisk. Alternatively, for a mapping with a 256 KB grain size, 8 KB of memory allows two mappings between one 8 GB source VDisk and two 8 GB target VDIs.

When you create a FlashCopy mapping, if you specify an I/O group other than the I/O group of the source VDisk, the memory accounting goes towards the specified I/O group, not towards the I/O group of the source VDisk.

An invocation example

```
svctask chiogrp -name testiogrhone io_grp0
```

The resulting output

No feedback

An invocation example for changing the amount of FlashCopy memory in io_grp0 to 30 MB

```
svctask chiogrp -feature flash -size 30 io_grp0
```

The resulting output

No feedback

chnode

You can use the **chnode** command to change the name that is assigned to a node and other options. The name can then be used when running subsequent commands. All the parameters that are associated with this command are optional. However, you must specify one or more parameters with this command.

Syntax

```
svctask -- chnode [ -iscsialias alias ] [ -noiscsialias ] [ -failover ] [ -name new_node_name ] [ node_name | node_id ]
```

Parameters

-iscsialias *alias*

(Optional) Specifies the iSCSI name of the node. The maximum length is 79 characters.

-noiscsialias

(Optional) Clears any previously set iSCSI name for this node. This parameter cannot be specified with the **iscsialias** parameter.

-failover

(Optional) Specifies that the name or iSCSI alias being set is the name or alias of the partner node in the I/O group. When there is no partner node, the values set are applied to the partner node when it is added to the cluster. If this parameter is used when there is a partner node, the name or alias of that node changes.

-name *new_node_name*

(Optional) Specifies the name to assign to the node.

| **Note:** Node names supplied with **-name** on **addnode** and **chnode** must not already be in use as node
| names or as node failover_names.

Important: The iSCSI Qualified Name (IQN) for each node is generated using the cluster and node names. If you are using the iSCSI protocol, changing either name also changes the IQN of all of the nodes in the cluster and might require reconfiguration of all iSCSI-attached hosts.

node_name | *node_id*

(Required) Specifies the node to be modified. The variable that follows the parameter is either:

- The node name that you assigned when you added the node to the cluster.
- The node ID that is assigned to the node (not the worldwide node name).

Description

If the **failover** parameter is not specified, this command changes the name or iSCSI alias of the node. The name can then be used to identify the node in subsequent commands.

The **failover** parameter is used to specify values that are normally applied to the partner node in the I/O group. When the partner node is offline, the iSCSI alias and IQN are assigned to the remaining node in the I/O Group. The iSCSI host data access is then preserved. If the partner node is offline when these parameters are set, the node they are set on handles iSCSI I/O requests to the iSCSI alias specified, or the IQN that is created using the node name. If the partner node in the I/O group is online when these parameters are set, the partner node handles iSCSI requests to the iSCSI alias specified, and its node name and IQN change.

An invocation example

```
svctask chnode -name testnodeone nodeone
```

The resulting output

No feedback

cleardumps

The **cleardumps** command cleans the various dump directories on a specified node.

Syntax

```
▶▶ svctask — —cleardumps— — -prefix — directory_or_file_filter — —————▶▶  
└─ node_id ─┬──  
└─ node_name ─┘
```

Parameters

-prefix *directory_or_file_filter*

(Required) Specifies the directory, files, or both to be cleaned. If a directory is specified, with no file filter, all relevant dump or log files in that directory are cleaned. You can use the following directory arguments (filters):

- **/dumps** (cleans all files in all subdirectories)
- **/dumps/cimom**
- **/dumps/configs**
- **/dumps/elogs**
- **/dumps/feature**
- **/dumps/iostats**
- **/dumps/iotrace**
- **/dumps/mdisk**
- **/home/admin/upgrade**

In addition to the directory, you can specify a filter file. For example, if you specify `/dumps/elogs/*.txt`, all files in the **/dumps/elogs** directory that end in `.txt` are cleaned.

Note: The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- With a wildcard, you must use double quotation marks (" ") around the filter entry, such as in the following entry:

```
>svctask cleardumps -prefix "/dumps/elogs/*.txt"
```

node_id | *node_name*

(Optional) Specifies the node to be cleaned. The variable that follows the parameter is either:

- The node name, that is, the label that you assigned when you added the node to the cluster

- The node ID that is assigned to the node (not the worldwide node name).

Description

This command deletes all the files that match the `directory/file_filter` argument on the specified node. If no node is specified, the configuration node is cleaned.

You can clean all the dumps directories by specifying `/dumps` as the directory variable.

You can clean all the files in a single directory by specifying one of the directory variables.

You can list the contents of these directories on the given node by using the `svcinfolsxdddumps` commands.

You can use this command to clean specific files in a given directory by specifying a directory or file name. You can use the wildcard character as part of the file name.

Note: To preserve the configuration and trace files, any files that match the following wildcard patterns are not cleaned:

- `*svc.config*`
- `*.trc`
- `*.trc.old`

An invocation example

```
svctask cleardumps -prefix /dumps/configs
```

The resulting output

No feedback

cpdumps

The `cpdumps` command copies dump files from a nonconfiguration node onto the configuration node.

Note: In the rare event that the `/dumps` directory on the configuration node is full, the copy action ends when the directory is full and provides no indicator of a failure. Therefore, clear the `/dumps` directory after migrating data from the configuration node.

Syntax

```

▶▶— svctask — — cpdumps — — -prefix ———┬── directory ———┬── node_name ———▶▶
                                   │   │   │   │   │   │   │   │   │
                                   └───┬── file_filter ───┬── node_id ───┬──

```

Parameters

-prefix *directory* | *file_filter*

(Required) Specifies the directory, or files, or both to be retrieved. If a directory is specified with no file filter, all relevant dump or log files in that directory are retrieved. You can use the following directory arguments (filters):

- `/dumps` (retrieves all files in all subdirectories)
- `/dumps/audit`
- `/dumps/cimom`
- `/dumps/configs`
- `/dumps/elogs`

- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/mdisk
- /home/admin/upgrade
-

- (Storwize V7000) /dumps/enclosure

In addition to the directory, you can specify a file filter. For example, if you specified /dumps/eologs/*.txt, all files in the /dumps/eologs directory that end in .txt are copied.

Note: The following rules apply to the use of wildcards with the CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, you must surround the filter entry with double quotation marks (""), as follows:

```
>svctask clear.dumps -prefix "/dumps/eologs/*.txt"
```

node_id | *node_name*

(Required) Specifies the node from which to retrieve the dumps. The variable that follows the parameter can be one of the following:

- The node name, or label that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the worldwide node name).

If the node specified is the current configuration node, no file is copied.

Description

This command copies any dumps that match the directory or file criteria from the given node to the current configuration node.

You can retrieve dumps that were saved to an old configuration node. During failover processing from the old configuration node to another node, the dumps that were on the old configuration node are not automatically copied. Because access from the CLI is only provided to the configuration node, cluster files can only be copied from the configuration node. This command enables you to retrieve files and place them on the configuration node so that you can then copy them.

You can view the contents of the directories by using the svcinfo lsxxxxdumps commands.

An invocation example

```
svctask cp.dumps -prefix /dumps/configs nodeone
```

The resulting output

No feedback

detectmdisk

The **detectmdisk** command allows you to manually rescan the fibre-channel network for any new managed disks (MDisks) that might have been added and to rebalance MDisk access across all available controller device ports.

Syntax

▶— svctask — — detectmdisk —————▶

Description

This command causes the cluster to rescan the fibre-channel network. The rescan discovers any new MDisks that have been added to the cluster and rebalances MDisk access across the available controller device ports. This command also detects any loss of controller port availability, and updates the SAN Volume Controller configuration to reflect any changes.

Note: Although it might appear that the **detectmdisk** command has completed, some extra time might be required for it to run. The **detectmdisk** is asynchronous and returns a prompt while the command continues to run in the background. You can use the **lsdiscoverystatus** command to list the discovery status.

In general, the cluster automatically detects disks when they appear on the network. However, some fibre-channel controllers do not send the required SCSI primitives that are necessary to automatically discover the new disks.

If you have attached new storage and the cluster has not detected it, you might need to run this command before the cluster detects the new disks.

When back-end controllers are added to the fibre-channel SAN and are included in the same switch zone as a cluster, the cluster automatically discovers the back-end controller and determines what storage is presented to it. The SCSI LUs that are presented by the back-end controller are displayed as unmanaged MDisks. However, if the configuration of the back-end controller is modified after this has occurred, the cluster might be unaware of these configuration changes. Run this command to rescan the fibre-channel network and update the list of unmanaged MDisks.

Note: The automatic discovery that is performed by the cluster does not write to an unmanaged MDisk. Only when you add an MDisk to an MDisk group, or use an MDisk to create an image mode virtual disk, is the storage actually used.

To identify the available MDisks, issue the **svctask detectmdisk** command to scan the fibre-channel network for any MDisks. When the detection is complete, issue the **svcinfolsmdiskcandidate** command to show the unmanaged MDisks; these MDisks have not been assigned to an MDisk group. Alternatively, you can issue the **svcinfolsmdisk** command to view all of the MDisks.

If disk controller ports have been removed as part of a reconfiguration, the SAN Volume Controller detects this change and reports the following error because it cannot distinguish an intentional reconfiguration from a port failure:

```
1630 Number of device logins reduced
```

If the error persists and redundancy has been compromised, the following more serious error is reported:

```
1627 Insufficient redundancy in disk controller connectivity
```

You must issue the **svctask detectmdisk** command to force SAN Volume Controller to update its configuration and accept the changes to the controller ports.

Note: Only issue the **svctask detectmdisk** command when all of the disk controller ports are working and correctly configured in the controller and the SAN zoning. Failure to do this could result in errors not being reported.

An invocation example

svctask detectmdisk

The resulting output

No feedback

ping

The **ping** command can be used to diagnose IP configuration problems by checking whether the specified IP address is accessible from the configuration node.

Syntax

```
▶▶ svctask — — ping — [ ipv4_address ] —————▶▶  
                        [ ipv6_address ]
```

Parameters

ipv4_address | *ipv6_address*

(Required) Specifies the cluster IP address.

Description

This command checks whether the specified IP address is accessible from the configuration node. The ping takes place only from the configuration node. It can be useful for diagnosing problems where the configuration node cannot be reached from a specific management server.

An invocation example

```
svctask ping 9.20.136.11
```

The resulting output

```
PING 9.20.136.11 (9.20.136.11) 56(84) bytes of data.  
64 bytes from 9.20.136.11: icmp_seq=1 ttl=249 time=0.690 ms  
64 bytes from 9.20.136.11: icmp_seq=2 ttl=249 time=0.382 ms  
64 bytes from 9.20.136.11: icmp_seq=3 ttl=249 time=0.311 ms  
--- 9.20.136.11 ping statistics ---  
3 packets transmitted, 3 received, 0% packet loss, time 2001ms  
rtt min/avg/max/mdev = 0.311/0.461/0.690/0.164 ms
```

rmnode

The **rmnode** command deletes a node from the cluster. You can enter this command any time after a cluster has been created.

Syntax

```
▶▶ svctask — — rmnode — [ -force ] [ node_name ] —————▶▶  
                        [ node_id ]
```

Parameters

-force

(Optional) Overrides the checks that this command runs. The parameter overrides the following two checks:

- If the command results in VDisks going offline, the command fails unless the **force** parameter is used.

- If the command results in a loss of data because there is unwritten data in the write cache that is contained only within the node to be removed, the command fails unless the **force** parameter is used.

If you use the **force** parameter as a result of an error about VDisks going offline, you force the node removal and run the risk of losing data from the write cache. The **force** parameter should always be used with caution.

node_name | *node_id*

Specifies the node to be deleted. The value for this parameter can be one of the following:

- The node name that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the worldwide node name).

Description

This command removes a node from the cluster. This makes the node a candidate to be added back into this cluster or into another cluster. After the node is deleted, the other node in the I/O group enters write-through mode until another node is added back into the I/O group.

By default, the **rmnode** command flushes the cache on the specified node before the node is taken offline. In some circumstances, such as when the system is already degraded (for example, when both nodes in the I/O group are online and the virtual disks within the I/O group are degraded), the system ensures that data loss does not occur as a result of deleting the only node with the cache data.

The cache is flushed before the node is deleted to prevent data loss if a failure occurs on the other node in the I/O group.

To take the specified node offline immediately without flushing the cache or ensuring data loss does not occur, run the **rmnode** command with the **-force** parameter.

Prerequisites:

Before you issue the **rmnode** command, perform the following tasks and read the following Attention notices to avoid losing access to data:

1. Determine which virtual disks (VDisks) are still assigned to this I/O group by issuing the following command. The command requests a filtered view of the VDisks, where the filter attribute is the I/O group.

```
svcinfolsvdisk -filtervalue IO_group_name=name
```

where *name* is the name of the I/O group.

Note: Any VDisks that are assigned to the I/O group that this node belongs to are assigned to the other node in the I/O group; the preferred node is changed. You cannot change this setting back.

2. Determine the hosts that the VDisks are mapped to by issuing the **svcinfolsvdiskhostmap** command.
3. Determine if any of the VDisks that are assigned to this I/O group contain data that you need to access:
 - If you *do not* want to maintain access to these VDisks, go to step 5 on page 125.
 - If you *do* want to maintain access to some or all of the VDisks, back up the data or migrate the data to a different (online) I/O group.
4. Determine if you need to turn the power off to the node:
 - If this is the last node in the cluster, you do not need to turn the power off to the node. Go to step 5 on page 125.

- If this is *not* the last node in the cluster, turn the power off to the node that you intend to remove. This step ensures that the Subsystem Device Driver (SDD) does not rediscover the paths that are manually removed before you issue the delete node request.
5. Update the SDD configuration for each virtual path (vpath) that is presented by the VDisks that you intend to remove. Updating the SDD configuration removes the vpaths from the VDisks. Failure to update the configuration can result in data corruption. See the *Multipath Subsystem Device Driver: User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.
 6. Quiesce all I/O operations that are destined for the node that you are deleting. Failure to quiesce the operations can result in failed I/O operations being reported to your host operating systems.

Attention:

1. Removing the last node in the cluster destroys the cluster. Before you delete the last node in the cluster, ensure that you want to destroy the cluster.
2. If you are removing a single node and the remaining node in the I/O group is online, the data can be exposed to a single point of failure if the remaining node fails.
3. This command might take some time to complete since the cache in the I/O group for that node is flushed before the node is removed. If the **-force** parameter is used, the cache is not flushed and the command completes more quickly. However, if the deleted node is the last node in the I/O group, using the **-force** option results in the write cache for that node being discarded rather than flushed, and data loss can occur. The **-force** option should be used with caution.
4. If both nodes in the I/O group are online and the VDisks are already degraded before deleting the node, redundancy to the VDisks is already degraded and loss of access to data and loss of data might occur if the **-force** option is used.

Notes:

1. If you are removing the configuration node, the **rmnode** command causes the configuration node to move to a different node within the cluster. This process might take a short time: typically less than a minute. The cluster IP address remains unchanged, but any SSH client attached to the configuration node might need to reestablish a connection. The management GUI reattaches to the new configuration node transparently.
2. If this is the last node in the cluster or if it is currently assigned as the configuration node, all connections to the cluster are lost. The user interface and any open CLI sessions are lost if the last node in the cluster is deleted. A time-out might occur if a command cannot be completed before the node is deleted.

An invocation example

```
svctask rmnode 1
```

The resulting output

```
No feedback
```

rmportip

The **rmportip** command removes an iSCSI IP address from a node ethernet port.

Syntax

```
| ►► svctask — — rmportip — [ -failover ] [ -ip_6 ] -node [ node_name ] —————►
|
| ►-port_id—►—————►—————►—————►—————►—————►—————►—————►—————►—————►
|
```

Parameters

-failover

(Optional) Specifies that the failover IP address information be removed for the specified port.

| **-ip_6**

(Optional) Specifies that the IPv6 address be removed for the specified port. If this parameter is not used, the IPv4 address is removed by default.

-node *node_name* | *node_id*

(Required) Specifies the node with the ethernet port that the IP address is being removed from.

| *port_id*

(Required) Specifies which port (1 or 2) to apply changes to.

Description

This command removes an IPv4 or IPv6 address from an ethernet port of a node.

An invocation example for IPv4

```
svctask rmportip -node 1 1
```

The resulting output

No feedback

An invocation example for IPv6

```
svctask rmportip -node 1 -ip_6 2
```

The resulting output

No feedback

setclustertime

The **setclustertime** command allows you to set the time for the cluster.

Syntax

```
▶▶— svctask — — setclustertime — — -time — time_value —————▶▶
```

Parameters

-time *time_value*

(Required) Specifies the time to which the cluster must be set. This must be in the following format:

MMDDHHmmYYYY

Description

This command sets the time for the cluster.

An invocation example

```
svctask setclustertime -time 040509142003
```

The resulting output

No feedback

setpwdreset

Use the **setpwdreset** command to view and change the status of the password-reset feature for the display panel.

Syntax

```
svctask -- setpwdreset -- [-disable | -enable | -show]
```

Parameters

-disable

Disables the password-reset feature that is available through the front panel menu system.

-enable

Enables the password-reset feature that is available through the front panel menu system.

-show

Displays the status of the password-reset feature, which is either enabled or disabled.

Description

The front panel menu system provides an option to reset the cluster superuser password. This option resets the password to a random string that is displayed on the front panel. You can then use this password to access the system. You can change the password at the next login.

Issue the **svctask setpwdreset** command to view and change the status of the password-reset feature for the display panel. Passwords can consist of the following characters: A - Z, a - z, 0 - 9, and underscore (_). Make a careful note of the cluster superuser password, because without it, you cannot access the cluster.

This command allows you access in case the cluster superuser password is forgotten. If you leave this feature enabled, you should ensure adequate physical security to the cluster hardware.

You can view or change the status of this feature.

An invocation example

```
svctask setpwdreset -show
```

The resulting output

```
Password status: [1]
```

This output means that the password or reset feature that is available through the front panel menu system is enabled. If the password status is [0], this feature is disabled.

settimezone

Use the **settimezone** command to set the time zone for the cluster.

Syntax

```
svctask -- settimezone -- -timezone -- timezone_arg
```

Parameters

-timezone *timezone_arg*

Specifies the time zone to set for the cluster.

Description

This command sets the time zone for the cluster. Use the **-timezone** parameter to specify the numeric ID of the time zone that you want to set. Issue the **svcinfolstimezones** command to list the time-zones that are available on the cluster. A list of valid time-zones settings are displayed in a list.

The time zone that this command sets will be used when formatting the error log that is produced by issuing the following command:

```
svctask dumperrlog
```

Note: If you have changed the timezone, you must clear the error log dump directory before you can view the error log through the Web application.

Issue the **svcinfolshowtimezone** command to display the current time-zone settings for the cluster. The cluster ID and its associated time-zone are displayed. Issue the **svctask setclustertime** command to set the time for the cluster.

An invocation example

```
svctask setttimezone -timezone 5
```

The resulting output

No feedback

startstats

- | Use the **startstats** command to modify the interval at which per-node statistics for virtual disks (VDisks),
- | managed disks (MDisks), and nodes are collected.

Syntax

```
▶▶ svctask — — startstats — — -interval — time_in_minutes —————▶▶
```

Parameters

-interval *time_in_minutes*

Specifies the time in minutes. This is the time interval between the gathering of statistics, from 1 to 60 minutes in increments of 1 minute.

Description

- | Running the **startstats** command will reset the statistics timer to zero (0), and give it a new interval at
- | which to sample. Statistics are collected at the end of each sampling period as specified by the **-interval**
- | parameter. These statistics are written to a file, with a new file created at the end of each sampling
- | period. Separate files are created for MDisks, VDIsks and node statistics.

The files generated are written to the `/dumps/iostats` directory.

A maximum of 16 files are stored in the directory at any one time for each statistics file type, for example:

Nm_stats_nodepanelname_date_time
Nv_stats_nodepanelname_date_time
Nn_stats_nodepanelname_date_time

Statistics files are created for all time intervals. Before the 17th file for each type is created, the oldest file of that type is deleted.

These files can be listed by using the `svcinfolsiostatsdumps` command.

The following naming convention is used for these files:

stats_type_stats_nodepanelname_date_time

Where *stats_type* is Nm for MDisks, Nv for VDIsks, and Nn for node statistics. *nodepanelname* is the current configuration node panel name, *date* is in the format of *yymmdd*, and *time* is in the format of *hhmmss*.

The following is an example of an MDisk statistics file name:

Nm_stats_000229_031123_072426

The following is an example of a VDisk statistics file name:

Nv_stats_000229_031123_072426

The following is an example of a node statistics file name:

Nn_stats_000229_031123_072426

Statistics are collected for each MDisk and recorded in the *Nm_stats_nodepanelname_date_time* file, including the following statistical information:

- The number of SCSI read and write commands that are processed during the sample period
- The number of blocks of data that are read and written during the sample period
- Per MDisk, cumulative read and write external response times in milliseconds
- Per MDisk, cumulative read and write queued response times

Statistics are collected for each VDisk and recorded in the *Nv_stats_nodepanelname_date_time* file, including the following statistical information:

- The total number of processed SCSI read and write commands
- The total amount of read and written data
- Cumulative read and write response time in milliseconds
- Statistical information about the read/write cache usage
- Global Mirror statistics including latency

Statistics are collected for the node from which the statistics file originated and recorded in the *Nn_stats_nodepanelname_date_time* file, including the following statistical information:

- Usage figure for the node from which the statistic file was obtained
- The amount of data transferred to and received from each port on the node to other devices on the SAN
- Statistical information about communication to other nodes on the fabric

An invocation example

```
svctask startstats -interval 25
```

The resulting output

No feedback

stopcluster

The **stopcluster** command allows you to shut down a single node or the entire cluster in a controlled manner. When you issue this command, you are prompted with a confirmation of intent to process the command.

Syntax

```
svctask — — stopcluster — [ -force ] [ -node [ node_name ] [ node_id ]
```

Parameters

-force

(Optional) Specifies that the node that is being shut down is the last online node in a given I/O group. The **force** parameter also overrides the checks that this command runs. The parameter overrides the following two checks:

- If the command results in VDisks going offline, the command fails unless the **force** parameter is used.
- If the node being shut down is the last online node in the I/O group, the command fails unless the **force** parameter is used.

If you use the **force** parameter as a result of an error about VDisks going offline, you force the node shutdown even if it is the last online node in the I/O group. The **force** parameter should always be used with caution.

-node *node_name* | *node_id*

(Optional) Specifies the node that you want to shut down. You can specify one of the following values:

- The node name, or label that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the worldwide node name).

If you specify **-node *node_name* | *node_id***, only the specified node is shut down; otherwise, the entire cluster is shut down.

Description

When you enter this command with no parameters, the entire cluster is shutdown. All data is flushed to disk before the power is removed.

Attention: Ensure that you have stopped all FlashCopy, Metro or Global Mirror, or data migration operations before you attempt to shutdown a node or cluster. You must also ensure that all asynchronous deletion operations have completed prior to a shutdown operation.

When you enter this command with either a node ID or node name, the node in question is shut down. After the command completes, the remaining node in the I/O group enters write-through mode until the power to the node is returned and the node rejoins the cluster.

Attention: If you have to remove all cluster input power for more than a few minutes, you must shut down the cluster before you remove the power. If the input power is removed from the uninterruptible power supply units without first shutting down the cluster and the uninterruptible power supply units, the uninterruptible power supply units remain operational and eventually become drained of power.

When input power is restored to the uninterruptible power supply units, they begin to recharge but the nodes do not permit I/O activity on the virtual disks until the uninterruptible power supply is charged

enough to enable all the data on the nodes to be saved in the event of an unexpected power loss. This might take as long as two hours. Shutting down the cluster prior to removing input power to the uninterruptible power supply units prevents the draining of battery power and makes it possible for I/O activity to resume as soon as input power is restored.

Attention: Before shutting down a node or the cluster, quiesce all I/O operations that are destined for this node or cluster. Failure to quiesce can result in failed I/O operations being reported to your host operating systems.

Begin the process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks that are provided by the cluster. Perform the following steps to quiesce all I/O to the cluster:

1. If you are unsure which hosts are using the VDisks that are provided by the cluster, determine the hosts that the VDisks are mapped to, as follows:
 - a. List the hosts that this VDisk is mapped to by issuing the following command, substituting the *vdiskname* | *id* variable with the name or ID of the VDisk:

```
svcinfolsvdiskhostmap vdiskname | id
```
 - b. In the command output, located the host name or ID to determine which host the specified VDisk is mapped to. If no data is displayed, the VDisk is not mapped to any hosts.
2. Repeat the previous procedure for all VDisks.

Attention: If you are shutting down the entire cluster, you lose access to all VDisks that are being provided by this cluster.

When all I/O has been stopped, issue the **svctask stopcluster** to shut down a single node or the entire cluster in a controller manner. If you specify the node ID or node name, you can shut down a single node. After the command completes, the other node in the I/O group goes into write-through mode until the power to the node is returned and the node rejoins the cluster.

Attention: If this is the last node in an I/O group, you must specify the **-force** parameter. You will lose all access to the virtual disks in the I/O group. Before you enter this command, ensure that this is what you want to do.

If a shutdown command has been sent to the cluster and both cluster and uninterruptible power supply units have powered off, when input power is restored, press the power button on the uninterruptible power supply front panel to restart the uninterruptible power supply units.

Ensure that you have stopped all FlashCopy mappings and Metro or Global Mirror relationships. In addition, ensure that all data migration operations and forced deletions have completed before continuing. Entering *y* to the confirmation message processes the command. No feedback is then displayed. Entering anything other than *y* or *Y* results in the command not processing. No feedback is displayed.

Attention: If you are shutting down a single node and the other node in the I/O group is online, the cache on the partner node goes into the write-through mode and that you are exposed to a single point of failure if the partner node fails while the node is shut down. You also lose access to all VDisks being served by this I/O group.

An invocation example

```
svctask stopcluster
```

The resulting output You will be presented with the following warning:

```
Are you sure that you want to continue with the shut down?
```

stopstats (Deprecated)

The **stopstats** command has been deprecated. You can no longer disable statistics collection.

Chapter 8. Backup and restore commands

The following commands are used for backing up and restoring configuration information with the SAN Volume Controller.

backup

Use the **backup** command to back up your configuration. You can enter this command any time after a cluster has been created.

Syntax

```
svsconfig -- backup [-quiet] [-v on | off]
```

Parameters

-quiet

Suppresses standard output (STDOUT) messages from the console.

-v on | off

On means verbose messages are displayed. Off means normal messages (the default) are displayed.

Description

- | The **backup** command extracts configuration information from the cluster, allowing you to restore your configuration whenever necessary. The **backup** command produces **svc.config.backup.xml**, **svc.config.backup.sh**, and **svsconfig.backup.log** files and saves them in the **/tmp** directory. The **.xml** file contains the extracted configuration information; the **.sh** file contains a script of the commands used to determine the configuration information; and the **.log** file contains details about command usage.
- | **Note:** If a previous **svc.config.backup.xml** file exists in **/tmp**, it is archived as **svc.config.backup.bak**; only one archive file is stored in the **/tmp** directory.
- | The underscore character (**_**) prefix is reserved for backup and restore command usage; do not use the underscore character in any object names.

An invocation example

```
svsconfig backup
```


The resulting output

```
No feedback
```

clear

Use the **clear** command to erase files in the **/tmp** directory that were previously produced by other **svsconfig** commands. You can enter this command any time after a cluster has been created.

Syntax



Parameters

-all

Erases all configuration files.

-q | quiet

Suppresses console output (STDOUT).

-v on | off

Produces verbose output (on); the default is regular output (off).

Description

This command erases configuration files on the current config node.

You can use the **svconfig clear** command without the **-all** parameter to erase files of the form:

```
/tmp/svc.config*.sh  
/tmp/svc.config*.log
```

You can use the **svconfig clear** command with the **-all** parameter to erase files of the form:

```
/tmp/svc.config*.sh  
/tmp/svc.config*.log  
/tmp/svc.config*.xml  
/tmp/svc.config*.bak
```

An invocation example

```
svconfig clear -all
```


The resulting output

No feedback

help

Use the **help** command to obtain summary information about the syntax of the **svconfig** command. You can enter this command any time after a cluster has been created.

Syntax



Parameters

-ver

Returns the version number for the **svconfig** command.

(action) -h | -?

Provides command help: the possible values for (action) are backup, clear, and restore.

-h | -?

Provides general help.

Description

This command provides syntax help for `svcconfig`.

An invocation example

```
svcconfig -ver
svcconfig -?
svcconfig backup -h
```


The resulting output

The help text displays.

restore

The **restore** command uses the configuration files in the `/tmp` directory to restore the cluster to its previous configuration.

Syntax



Parameters

-f | force

Forces continued processing where possible.

-q | quiet

Suppresses console output (STDOUT).

-prepare

| Checks the current configuration against the information in `svc.config.backup.xml` on the
| configuration to be restored. Prepares commands for processing in `svc.config.restore.sh`, and
| produces a log of events in `svc.config.restore.prepare`.

-fmt | fmtdisk

| Includes the `-fmtdisk` option on all `mkvdisk` commands to be issued. Not allowed with the `-execute`
| parameter.

-execute

Runs the command script `svc.config.restore.sh`. Produces a log of events in
`svc.config.restore.execute.log`.

-v on | off

Produces verbose output (on); the default is regular output (off).

Description

| The **restore** command restores the target cluster configuration from the **svc.config.backup.xml** file in the configuration files directory. If neither the **-prepare** nor **-execute** option is specified, the command performs both phases in sequence, and only a single event log, **svc.config.restore.log**, is produced.

| The restore operation is also called T4 (Tier 4) Recovery, and should be used on a cluster that has just been started. It should not be used on a cluster that already has any nonautomatic objects configured, such as MDisk groups (storage pools) or VDisks (volumes).

| The restore operation is performed in two phases:

- | 1. Prepare
- | 2. Execute

| The command pauses for 8 minutes if any nodes are added during this process. You are informed of this at run-time.

The configuration files directory is **/tmp**.

An invocation example

```
| svcconfig restore -prepare -fmt  
| svcconfig restore -execute  
| svcconfig restore
```

The resulting output

No feedback

Chapter 9. Cluster diagnostic and service-aid commands

Cluster diagnostic and service-aid commands are designed to diagnose and find cluster problems.

The SAN Volume Controller enables you to perform service activity, such as problem determination and repair activities, with a limited set of command-line tools. When you are logged in under the administrator role, all command-line activities are permitted. When you are logged in under the service role, only those commands that are required for service are enabled. All of these commands apply under the service role.

applysoftware

The **applysoftware** command upgrades the cluster to a new level of software.

Syntax

```
| ►► svctask — — applysoftware — — [ -file — filename_arg ] [ -abort ] [ -force ]
```

Parameters

- | **-file** *filename_arg*
| (Required for performing an upgrade) Specifies the file name of the new software package to be
| applied.
- | **-abort**
| (Required for stopping an upgrade) Specifies that a stalled upgrade should be stopped, returning the
| cluster to the original software level.

| **Note:** The **force** parameter can be used with the **abort** parameter. If one or more nodes are offline,
| you must use the **force** parameter with the **abort** parameter.
- | **-force**
| (Optional) Specifies that the upgrade or abort should proceed even if there are non-redundant nodes
| in the cluster.

| **Note:** Using this option might result in a loss of access.

Description

This command starts the upgrade process of the cluster to a new level of SAN Volume Controller software. The **applysoftware** command applies a level of software to the node as a service action (Paced Upgrade) to upgrade the specific node, or as an automatic upgrade process that upgrades all of the nodes in the entire cluster.

The software package as specified by the file name must first be copied onto the current configuration node in the /home/admin/upgrade directory. You can use the PuTTY secure copy (scp) application to copy the file.

The command completes as soon as the upgrade process has successfully begun. The command fails and the upgrade package is deleted if:

- The given package fails an integrity check due to corruption.
- Any node in the cluster has a hardware type not supported by the new software.

- The new software level does not support upgrades from the currently installed software.
- The software level of a remote cluster is incompatible with the new software.
- There are any VDisks that are dependent on the status of a node.

Note: The **force** parameter can be used to override this if you are prepared to lose access to data during the upgrade. Before proceeding, use the **svcinfoldependentvdisks** command with the **node** parameter to list the node-dependent VDisks at the time the command is run. If the command returns an error, you must move your quorum disks to MDisks that are accessible through all nodes. Rerun the command until no errors are returned.

The actual upgrade completes asynchronously.

The **svcinfoldumps** command allows you to view the contents of the `/home/admin/upgrade` directory.

An invocation example

```
svctask applysoftware -file softwareupdate
```

The resulting output

No feedback

cherrstate

The **cherrstate** command marks an unfixed error as fixed. You can also use it to mark a fixed error as unfixed.

Syntax

```
svctask -- cherrstate -- -sequencenumber sequence_number -- [-unfix]
```

Parameters

-sequencenumber *sequence_number*

(Required) Specifies the error log sequence numbers to mark as fixed or as unfixed.

-unfix

(Optional) Specifies that the sequence numbers be marked as unfixed. Use this parameter when you have marked the wrong sequence number as fixed.

Description

The error log entries that the sequence numbers that you entered are marked as fixed. Use this command as a manual confirmation step that you have performed a maintenance procedure on the cluster, the fabric, or the subsystems.

This step is performed as part of the directed maintenance procedures (DMPs).

Optionally, if you have wrongly marked a sequence number as fixed, you can use the **-unfix** parameter to change the entry to unfixed.

An invocation example

```
svctask cherrstate -sequencenumber 2019
```

The resulting output

No feedback

clearerrlog

The **clearerrlog** command clears all entries from the error log including status events and any unfixed errors.

Syntax

```
svctask -- clearerrlog -- [-force]
```

Parameters

-force

(Optional) Specifies that the **clearerrlog** command be processed without confirmation requests. If the **-force** parameter is not supplied, you are prompted to confirm that you want to clear the log.

Description

This command clears all entries from the error log. The entries are cleared even if there are unfixed errors in the log. It also clears any status events that are in the log.

Attention: This command is destructive. Use it only use when you have either rebuilt the cluster or have fixed a major problem that has caused entries in the error log that you do not want to manually fix.

An invocation example

```
svctask clearerrlog -force
```

The resulting output

No feedback

dumperrlog

The **dumperrlog** command dumps the contents of the error log to a text file.

Syntax

```
svctask -- dumperrlog -- [-prefix filename_prefix]
```

Parameters

-prefix *filename_prefix*

(Optional) A file name is created from the prefix and a time stamp, and has the following format:

```
prefix_NNNNNN_YYMMDD_HHMMSS
```

where *NNNNNN* is the node front panel name.

Note: If the **-prefix** parameter is not supplied, the dump is directed to a file with a system-defined prefix of **errlog**.

Description

When run with no parameters, this command dumps the cluster error log to a file using a system-supplied prefix of `errlog`, which includes the node ID and time stamp. When a file name prefix is provided, the same operation is performed but the details are stored in the `dumps` directory within a file with a name that starts with the specified prefix.

A maximum of ten error-log dump files are kept on the cluster. When the 11th dump is made, the oldest existing dump file is overwritten.

Error log dump files are written to `/dumps/elogs`. The contents of this directory can be viewed using the `svcinfolerrlogdumps` command.

Files are not deleted from other nodes until you issue the `cleardumps` command.

An invocation example

```
svctask dumperrlog -prefix testerrorlog
```

The resulting output

No feedback

finderr

The `finderr` command analyzes the error log for the highest severity unfixed error.

Syntax

```
svctask — — finderr —————>
```

Description

The command scans the error log for any unfixed errors. Given a priority ordering within the code, the highest priority unfixed error is returned to standard output.

You can use this command to determine the order in which to fix the logged errors.

An invocation example

```
svctask finderr
```

The resulting output

Highest priority unfixed error code is [1010]

setlocale

The `setlocale` command changes the locale setting for the cluster. It also changes command output to the chosen language.

Syntax

```
svctask — — setlocale — — -locale — locale_id —————>
```

Parameters

-locale *locale_id*
Specifies the locale ID.

Description

This command changes the language in which error messages are displayed as output from the command-line interface. Subsequently, all error messages from the command-line tools are generated in the chosen language. This command is run when you request a change of language (locale) and is generally run from the Web page. Issue the **setlocale** command to change the locale setting for the cluster; all interface output is changed to the chosen language. For example, to change the language to Japanese, type the following:

```
svctask setlocale -locale 3
```

where **3** is the value for Japanese. The following values are supported:

- **0** US English (default)
- | • **1** Simplified Chinese
- | • **2** Traditional Chinese
- **3** Japanese
- | • **4** French
- | • **5** German
- | • **6** Italian
- | • **7** Spanish

Note: This command does not change the front panel display panel settings.

An invocation example (where **3** is Japanese)

```
svctask setlocale -locale 3
```

The resulting output

No feedback

svqueryclock

The **svqueryclock** command returns the date, time, and current time-zone of the cluster.

Syntax

▶▶—svqueryclock—▶▶

Description

This command returns the date, time and current time-zone of the cluster.

An invocation example

```
svqueryclock
```

The resulting output

```
Mon Nov 25 14:59:28 GMT 2002
```

writesernum

Use the **writesernum** command to write the node serial number into the planar NVRAM.

Syntax

```
svctask writesernum -senum serial_number [node_id | node_name]
```

Parameters

-senum *serial_number*

(Required) Specifies the serial number to write to the nonvolatile memory of the system planar.

node_id | *node_name*

(Required) Specifies the node where the system planar is located. The serial number is written to this system planar. This name is not the worldwide node name (WWNN).

Description

This command writes the node serial number into the planar NVRAM and then reboots the system. You can find the serial number at the front of the node without having to remove it from the rack. The seven-digit alphanumeric serial number is located on a label on the front of the node. The serial number on the label might contain a hyphen. Omit this hyphen when typing the serial number with the **writesernum** command.

Note: Once you have written the serial number to the planar NVRAM, you can issue the **svcinfo lsnodevpd** command to verify that the number is correct. The `system_serial_number` field contains the serial number.

An invocation example

```
svctask writesernum -senum 1300027 node1
```

The resulting output

No feedback

Chapter 10. Host commands

The following commands enable you to work with host objects with the SAN Volume Controller.

addhostiogrp

The **addhostiogrp** command enables you to map I/O groups to an existing host object.

Syntax

```
svctask -- addhostiogrp -- [-iogrp -- iogrp_list] [-iogrpall] [ host_name | host_id ]
```

Parameters

-iogrp *iogrp_list*

(Required if you do not use **-iogrpall**) Specifies a colon-separated list of one or more I/O groups that must be mapped to the host. You cannot use this parameter with the **-iogrpall** parameter.

-iogrpall

(Required if you do not use **-iogrp**) Specifies that all the I/O groups must be mapped to the specified host. You cannot use this parameter with the **-iogrp** parameter.

host_id | *host_name*

(Required) Specifies the host to which the I/O groups must be mapped, either by ID or by name.

Description

This command allows you to map the list of I/O groups to the specified host object.

An invocation example

```
svctask addhostiogrp -iogrpall testhost
```

The resulting output

No feedback

addhostport

The **addhostport** command adds worldwide port names (WWPNs) or iSCSI names to an existing host object.

Syntax

```
svctask -- addhostport -- [-hbawwpn -- wwpn_list] [-iscsiname--iscsi_name_list] [-force] [ host_name | host_id ]
```

Parameters

-hbawwpn *wwpn_list*

(Required if you do not use **iscsiname**) Specifies the list of fibre-channel host ports to add to the host. At least one WWPN or iSCSI name must be specified. You cannot use this parameter with the **iscsiname** parameter.

-iscsiname *iscsi_name_list*

(Required if you do not use **hbawwpn**) Specifies the comma-separated list of iSCSI names to add to the host. At least one WWPN or iSCSI name must be specified. You cannot use this parameter with the **hbawwpn** parameter.

-force

(Optional) Specifies that the list of ports be added to the host without the validation of any WWPNs or iSCSI names.

host_id | *host_name*

(Required) Specifies the host object to add ports to, either by ID or by name.

Description

This command adds a list of HBA WWPNs or iSCSI names to the specified host object. Any virtual disks that are mapped to this host object automatically map to the new ports.

Only WWPNs that are logged-in unconfigured can be added. For a list of candidate WWPNs, use the **svcinfolshbaportcandidate** command.

Some HBA device drivers do not log in to the fabric until they can recognize target LUNs. Because they do not log in, their WWPNs are not recognized as candidate ports. You can specify the **force** parameter with the **svctask addhostport** command to stop the validation of the WWPN list.

Note: When all I/O groups are removed from an iSCSI host, you cannot add a port to the iSCSI host until you have mapped the iSCSI host to at least one I/O group. After mapping the iSCSI host to at least one I/O group, resubmit the **svctask addhostport** command. After adding the port to the host, you must create a host authentication entry using the **svctask chhost** command.

An invocation example

```
svctask addhostport -hbawwpn 210100E08B251DD4 host_one
```

The resulting output

No feedback

An invocation example

```
svctask addhostport -iscsiname iqn.localhost.hostid.7f000001 mhost13
```


The resulting output

No feedback

chhost

The **chhost** command changes the name or type of a host object. This does not affect any existing virtual disk-to-host mappings.

Syntax



Parameters

-type *hpux* | *tpgs* | *generic*

(Optional) Specifies the type of host: **hpux**, **tpgs**, or **generic**. The default is **generic**. The **tpgs** parameter enables extra target port unit attentions. Refer to SAN Volume Controller host attachment documentation for more information on the hosts that require the **type** parameter.

-name *new_name_arg*

(Optional) Specifies the new name that you want to assign to the host object.

-mask *port_login_mask*

(Optional) Specifies which node target ports that a host can access. The port mask is four binary bits and is made of a combination of 0's and 1's, where 0 indicates that the corresponding target port cannot be used and 1 indicates that it can be used. The right-most bit in the mask corresponds to the lowest numbered target port (1 not 4) on a node. Valid mask values range from **0000** (no ports enabled) to **1111** (all ports enabled). For example, a mask of **0011** enables port 1 and port 2. The default value is **1111** (all ports enabled).

-chapsecret *chap_secret*

(Optional) Sets the Challenge Handshake Authentication Protocol (CHAP) secret used to authenticate the host for iSCSI I/O. This secret is shared between the host and the cluster. The CHAP secret for each host can be listed using the **svcinfolsiscsi** command.

-nochapsecret

(Optional) Clears any previously set CHAP secret for this host.

host_name | *host_id*

(Required) Specifies the host object to modify, either by ID or by current name.

Description

This command can change the name of the specified host to a new name, or it can change the type of host. This command does not affect any of the current virtual disk-to-host mappings.

The port mask applies to logins from the host initiator port that are associated with the host object. For each login between a host HBA port and node port, the node examines the port mask that is associated with the host object for which the host HBA is a member and determines if access is allowed or denied. If access is denied, the node responds to SCSI commands as if the HBA port is unknown.

Note: When all I/O groups are removed from an iSCSI host, the **svcinfolsiscsi** command does not display the authentication entry for that host. Use the **svctask addhostiogr** command to map the iSCSI host to at least one I/O group, and then use the **svctask addhostport** command to add the iSCSI port into it. You must also add authentication for that host using the **svctask chhost** command with either the **chapsecret** or **nochapsecret** parameter.

An invocation example

```
svctask chhost -name testhostlode -mask 0011 hostone
```

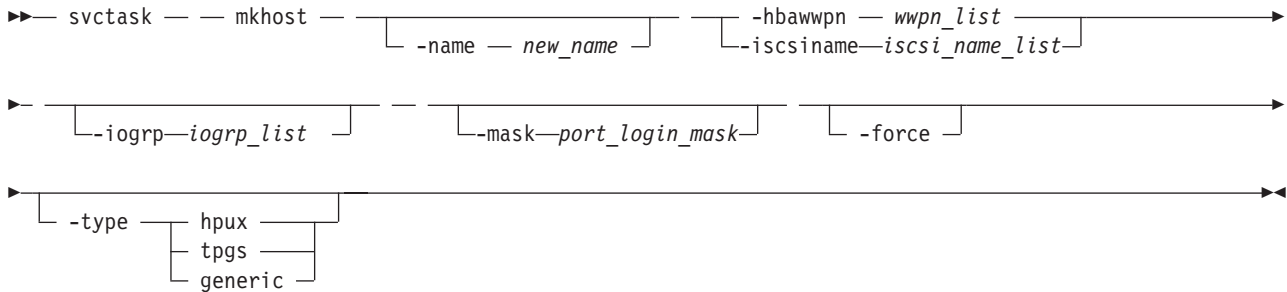
The resulting output

No feedback

mkhost

The **mkhost** command creates a logical host object.

Syntax



Parameters

- name** *new_name*
(Optional) Specifies a name or label for the new host object.
- hbawwpn** *wwpn_list*
(Required if you do not use **iscsiname**) Specifies one or more host bus adapter (HBA) worldwide port names (WWPNs) to add to the specified host object. At least one WWPN or iSCSI name must be specified. You cannot use this parameter with the **iscsiname** parameter.
- iscsiname** *iscsi_name_list*
(Required if you do not use **hbawwpn**) Specifies one or more iSCSI names of this host. At least one WWPN or iSCSI name must be specified. You cannot use this parameter with the **hbawwpn** parameter.
- iogrp** *iogrp_list*
(Optional) Specifies a set of one or more I/O groups that the host can access the VDisks from. I/O groups are specified using their names or IDs, separated by a colon. Names and IDs can be mixed in the list. If this parameter is not specified, the host is associated with all I/O groups.
- mask** *port_login_mask*
(Optional) Specifies which node target ports that a host can access. The port mask is four binary bits and is made up of a combination of 0's and 1's, where 0 indicates that the corresponding target port cannot be used and 1 indicates that it can be used.. The right-most bit in the mask corresponds to the lowest numbered target port (1 not 4) on a node. Valid mask values range from **0000** (no ports enabled) to **1111** (all ports enabled). For example, a mask of **0011** enables port 1 and port 2. The default value is **1111** (all ports enabled).
- force**
(Optional) Specifies that a logical host object be created without validation of the WWPNs.
- type** **hpux** | **tpgs** | **generic**
(Optional) Specifies the type of host: **hpux**, **tpgs**, or **generic**. The default is **generic**. The **tpgs** parameter enables extra target port unit attentions. Refer to SAN Volume Controller host attachment documentation for more information on the hosts that require the **type** parameter.

Description

The **mkhost** command associates one or more HBA WWPNs or iSCSI names with a logical host object. This command creates a new host. The ID is displayed when the command completes. You can subsequently use this object when you map virtual disks to hosts by using the **mkvdiskhostmap** command.

Issue the **mkhost** command only once. The cluster scans the fabric for WWPNs in the host zone. The cluster itself cannot filter into the hosts to determine which WWPNs are in which hosts. Therefore, you must use the **svctask mkhost** command to identify the hosts.

After you identify the hosts, mappings are created between hosts and virtual disks. These mappings effectively present the virtual disks to the hosts to which they are mapped. All WWPNs in the host object are mapped to the virtual disks.

Some HBA device drivers do not log in to the fabric until they can see target logical unit numbers (LUNs). Because they do not log in, their WWPNs are not recognized as candidate ports. You can specify the **force** parameter with this command to stop the validation of the WWPN list.

This command fails if you add the host to an I/O group that is associated with more host ports or host objects than is allowed by the limits within the cluster.

For additional information, see the **svctask mkvdiskhostmap** and **svcinfo lshbaportcandidate** commands.

An invocation example

```
svctask mkhost -name hostone -hbawwpn 210100E08B251DD4 -force -mask 1001
```

The resulting output

```
Host id [1] successfully created
```

An invocation example

```
svctask mkhost -iscsiname iqn.localhost.hostid.7f000001 -name newhost
```

The resulting output

```
Host, id [10], successfully created
```

rmhost

The **rmhost** command deletes a host object.

Syntax

```
➤— svctask — — rmhost — [ -force ] [ host_name | host_id ] —————➤
```

Parameters

-force

(Optional) Specifies that you want the system to delete the host object even if mappings still exist between this host and virtual disks (VDisks). When the **-force** parameter is specified, the mappings are deleted before the host object is deleted.

host_name | *host_id*

(Required) Specifies the host object to delete, either by ID or by name.

Description

The **rmhost** command deletes the logical host object. The WWPNs that were contained by this host object (if it is still connected and logged in to the fabric) are returned to the unconfigured state. When you issue the **svcinfo lshbaportcandidate** command, the host objects are listed as candidate ports.

If any mappings still exist between this host and virtual disks, the command fails unless you specify the **-force** parameter. When the **-force** parameter is specified, the **rmhost** command deletes the mappings before the host object is deleted.

An invocation example

```
svctask rmhost host_one
```

The resulting output

No feedback

rmhostiogr

The **rmhostiogr** command enables you to delete mappings between one or more I/O groups and a specified host object.

Syntax

```
svctask -- rmhostiogr -- [-iogrp -- iogrp_list] [-iogrpall] [-force] [host_name | host_id]
```

Parameters

-iogrp *iogrp_list*

(Required) Specifies a set of one or more I/O group mappings that will be deleted from the host. You cannot use this parameter with the **-iogrpall** parameter.

-iogrpall

(Optional) Specifies that all the I/O group mappings that are associated with the specified host must be deleted from the host. You cannot use this parameter with the **-iogrp** parameter.

-force

(Optional) Specifies that you want the system to remove the specified I/O group mappings on the host even if the removal of a host to I/O group mapping results in the loss of VDisk-to-host mappings.

host_id | *host_name*

(Required) Specifies the identity of the host either by ID or name from which the I/O group mappings must be deleted.

Description

The **rmhostiogr** command deletes the mappings between the list of I/O groups and the specified host object.

This command fails if any of the I/O groups that are deleted from the host contain VDIsks that have host mappings to the host. To resolve this problem, do one of the following:

- Delete the VDIsks-to-host mappings that are causing the error.
- Delete the VDIsks or the host.
- Issue the **rmhostiogr** command with the **-force** parameter.

Note: When all I/O groups are removed from an iSCSI host, and you want to add an iSCSI port to the host, refer to the **svctask addhostport** and **svctask chhost** commands.

An invocation example

```
svctask rmhostiogr -iogrp 1:2 host0
```

The resulting output

No feedback

rmhostport

The **rmhostport** command deletes worldwide port names (WWPNs) or iSCSI names from an existing host object.

Syntax

```
►►— svctask — — rmhostport — — [ -hbawwpn — wwpn_list ] [ -iscsiname — iscsi_name_list ] [ -force ] [ host_name | host_id ] ►►
```

Parameters

-hbawwpn *wwpn_list*

(Required if you do not use **iscsiname**) Specifies the list of fibre-channel host ports to delete from the host. At least one WWPN or iSCSI name must be specified. You cannot use this parameter with the **iscsiname** parameter.

-iscsiname *iscsi_name_list*

(Required if you do not use **hbawwpn**) Specifies the comma-separated list of iSCSI names to delete from the host. At least one WWPN or iSCSI name must be specified. You cannot use this parameter with the **hbawwpn** parameter.

-force

(Optional) Forces the deletion of the specified ports. This overrides the check that all of the WWPNs or iSCSI names in the list are mapped to the host specified.

host_name | *host_id*

(Required) Specifies the host name or the host ID.

Description

This command deletes the list of HBA WWPNs or iSCSI names from the specified host object. If the WWPN ports are still logged in to the fabric, they become unconfigured and are listed as candidate WWPNs. See also the **svcinfolshbaportcandidate** command.

Any virtual disks that are mapped to this host object are automatically unmapped from the ports.

Replacing an HBA in a host: List the candidate HBA ports by issuing the **svcinfolshbaportcandidate** command. A list of the HBA ports that are available to be added to host objects is displayed. One or more of these ports corresponds with one or more WWPNs that belong to the new HBA. Locate the host object that corresponds to the host in which you have replaced the HBA. The following command lists all the defined host objects:

```
svcinfolshost
```

To list the WWPNs that are currently assigned to the host, issue the following:

```
svcinfolshost hostobjectname
```

where *hostobjectname* is the name of the host object.

Add the new ports to the existing host object by issuing the following command:

```
svctask addhostport -hbawwpn one or more existing WWPNs  
separated by : hostobjectname/ID
```

where *one or more existing WWPNs separated by :* and *hostobjectname/id* correspond to those values listed in the previous steps.

Remove the old ports from the host object by issuing the following command:

```
svctask rmhostport -hbawwpn one or more existing WWPNs  
separated by : hostobjectname/ID
```

where *one or more existing WWPNs separated by :* corresponds with those WWPNs that are listed in the previous step that belong to the old HBA that has been replaced. Any mappings that exist between the host object and VDisks are automatically applied to the new WWPNs. Therefore, the host recognizes that the VDisks are the same SCSI LUNs as before. See the *Multipath Subsystem Device Driver: User's Guide* for additional information about dynamic reconfiguration.

An invocation example

```
svctask rmhostport -hbawwpn 210100E08B251DD4 host_one
```

The resulting output

No feedback

An invocation example

```
svctask rmhostport -iscsiname iqn.localhost.hostid.7f000001 mchost13
```

The resulting output

No feedback

Chapter 11. Virtual disk commands

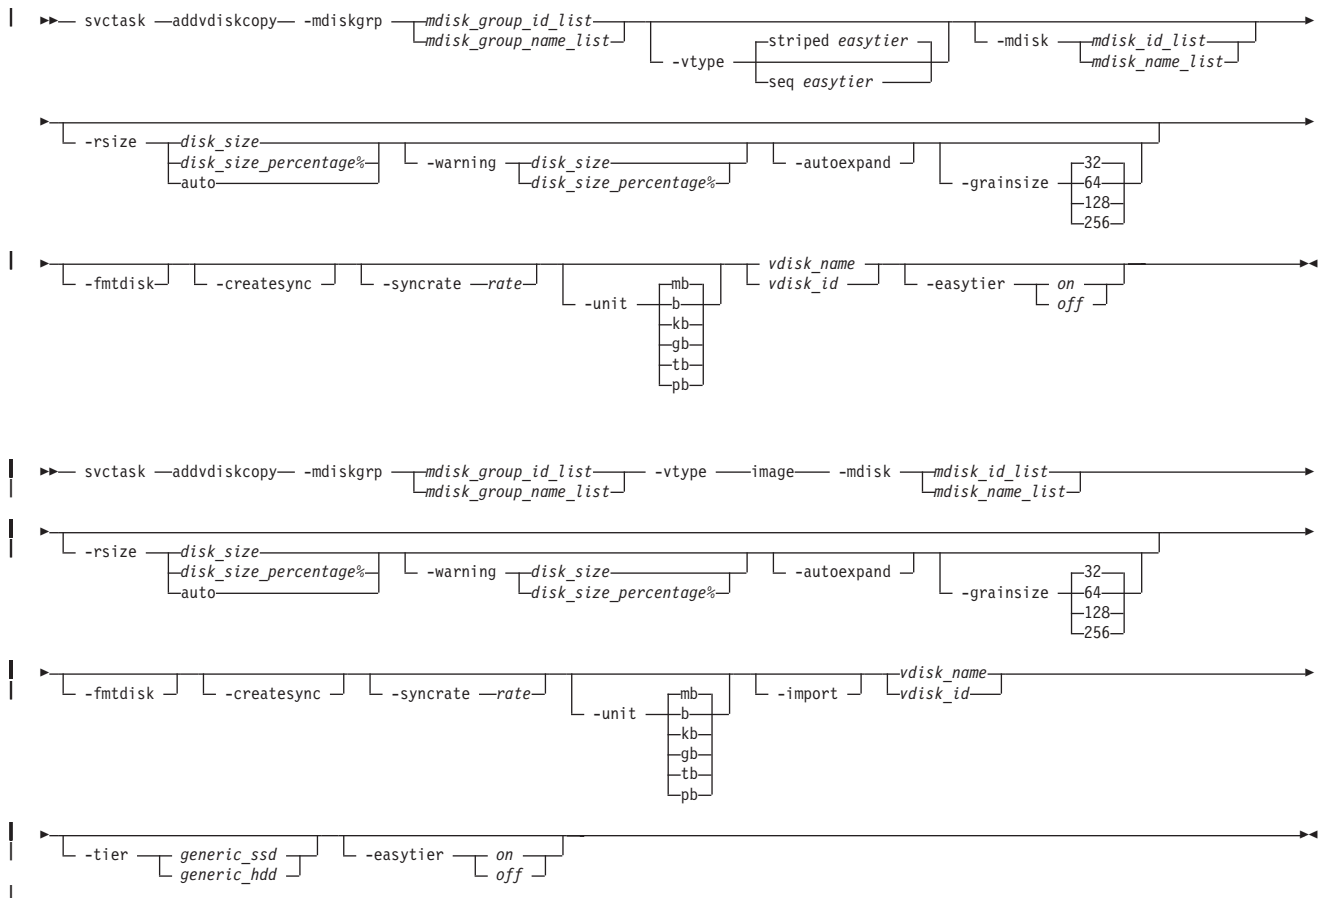
The following commands enable you to work with virtual disk options with the SAN Volume Controller.

addvdiskcopy

The **addvdiskcopy** command adds a copy to an existing VDisk, which changes a nonmirrored VDisk into a mirrored VDisk.

Note: The first syntax diagram depicts the addition of a sequential or striped mode virtual disk. The second syntax diagram depicts the addition of an image mode virtual disk.

Syntax



Parameters

-mdiskgrp *mdisk_group_id_list* | *mdisk_group_name_list*

(Required) Specifies the managed disk groups to use to create copies for the virtual disk. You must specify a group for each copy that is being added.

-vtype *seq* | *striped* | *image*

(Optional) Specifies the virtualization type for the copy: sequential, striped, or image. The type can be different than the virtualization types for other copies on the VDisk. The default virtualization type is **striped**.

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Optional) Specifies one or more managed disks (MDisks). For sequential and image mode copies, you must specify a single MDisk that has sufficient free extents. For image mode copies, the MDisk must be in unmanaged mode. For sequential mode copies the MDisk must be in the managed mode.

-syncrate *rate*

(Optional) Specifies the copy synchronization rate. A value of zero (0) prevents synchronization. The default value is 50. For the supported **-syncrate** values and their corresponding rates, see Table 12 on page 154.

-createsync

(Optional) Suppresses the synchronization of the new VDisk copy with the primary copy. Using this parameter can cause data corruption if the primary copy fails and leaves an unsynchronized secondary copy to provide data. Using this parameter can cause loss of read stability in unwritten areas if the primary copy fails, data is read from the primary copy, and then different data is read from the secondary copy. To avoid data loss or read stability loss, use this parameter only for a primary copy that has been formatted and not written to, and with the **-fmtdisk** parameter.

-fmtdisk

(Optional) Formats a sequential or striped mode copy. You must also specify the **-createsync** parameter, which labels the formatted copy as identical to the primary copy. The **-fmtdisk** parameter causes the VDisk to go offline until new VDisk copy formatting completes. To query the formatting progress, use the **lsvdiskprogress** command.

-rsize *disk_size* | *disk_size_percentage%* | **auto**

(Optional) Makes the copy space-efficient and specifies the real size of the copy. Specify the *disk_size* | *disk_size_percentage* value using an integer, or an integer immediately followed by the percent character (%). The default units for *disk_size* are megabytes (MB); to specify different units, use the **-unit** parameter. The **auto** option creates a VDisk copy that uses the entire size of the MDisk; if you specify the **-rsize auto** option, you must also specify the **-vtype image** option.

-warning *disk_size* | *disk_size_percentage%*

(Optional) Requires that the **-rsize** parameter also be specified. Generates a warning when the used disk capacity on the space-efficient copy first exceeds the specified threshold. You can specify a *disk_size* integer, which defaults to megabytes (MB) unless the **-unit** parameter is specified; or you can specify a *disk_size%*, which is a percentage of the virtual disk size. If **-autoexpand** is enabled, the default value for **-warning** is 80% of the virtual disk capacity. If **-autoexpand** is not enabled, the default value for warning is 80% of the real capacity. To disable warnings, specify 0.

-autoexpand

(Optional) Requires that the **-rsize** parameter also be specified. Specifies that space-efficient copies automatically expand their real capacities by allocating new extents from their managed disk group. If the **-autoexpand** parameter is specified, the **-rsize** parameter specifies a capacity that is reserved by the copy. This protects the copy from going offline when its managed disk group runs out of space by allowing it to consume this reserved space first.

-grainsize 32 | 64 | 128 | 256

(Optional) Requires that the **-rsize** parameter also be specified. Sets the grain size (KB) for a space-efficient VDisk copy. The default is 32 KB.

-unit b | kb | mb | gb | tb | pb

(Optional) Specifies the data units for the **-rsize** and **-warning** parameters.

-import

(Optional) Imports an image mode disk that contains a space-efficient volume into the cluster. Requires that the **-rsize** and **-vtype image** parameters also be specified.

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk to add the VDisk copy to, either by ID or by name.

| **-tier**
 | (Optional) Specifies the MDisk tier when an image mode copy is added.

| **-easytier**
 | (Optional) Determines if the IBM System Storage Easy Tier function is allowed to move extents for this VDisk. If a volume copy is striped and not being migrated the following table applies:

| *Table 11. Storage pool Easy Tier settings*

Storage pool Easy Tier setting	Number of tiers in the storage pool	Volume copy Easy Tier setting	Volume copy Easy Tier status
Off	One	Off	inactive (see note 2)
Off	One	On	inactive (see note 2)
Off	Two	Off	inactive (see note 2)
Off	Two	On	inactive (see note 2)
Auto (see note 5)	One	Off	inactive (see note 2)
Auto (see note 5)	One	On	inactive (see note 2)
Auto (see note 5)	Two	Off	measured (see note 3)
Auto (see note 5)	Two	On	active (see note 1)
On (see note 5)	One	Off	measured (see note 3)
On (see note 5)	One	On	measured (see note 3)
On (see note 5)	Two	Off	measured (see note 3)
On (see note 5)	Two	On	active (see note 1)

Note:

1. If the volume copy is in image or sequential mode or is being migrated then the volume copy Easy Tier status will be *measured* instead of *active*.
2. When the volume copy status is *inactive*, no Easy Tier functions are enabled for that volume copy.
3. When the volume copy status is *measured*, the Easy Tier function collects usage statistics for the volume but automatic data placement is not active.
4. When the volume copy status is *active*, the Easy Tier function operates in automatic data placement mode for that volume.
5. The default Easy Tier setting for a storage pool is *auto*, and the default Easy Tier setting for a volume copy is *on*. This means that Easy Tier functions will be disabled for storage pools with a single tier, and that automatic data placement mode will be enabled for all striped volume copies in a storage pool with two tiers.

| Description

The **addvdiskcopy** command adds a copy to an existing VDisk, which changes a nonmirrored VDisk into a mirrored VDisk. Use the **-mdiskgrp** parameter to specify the managed disk group that will provide storage for the copy; the **svcinfo lsmdiskgrp** command lists the available managed disk groups and the amount of available storage in each group.

The virtualization types are defined as follows:

sequential (seq)

This policy requires the **-mdisk** parameter with a single managed disk as its argument. This MDisk must be in the managed mode.

It creates the virtual disk using extents from the given managed disk (assuming there are enough free extents on the managed disk).

striped

This is the default policy. If the **-vtype** parameter is not specified, this policy is used in its default

form. That is, all managed disks in the managed disk group are used to create the virtual disk. The striping is at an extent level; one extent from each managed disk in the group is used. For example, a managed disk group with 10 managed disks uses one extent from each managed disk, then it uses the 11th extent from the first managed disk, and so on.

If the **-mdisk** parameter is also specified, you can supply a list of managed disks to use as the stripe set. This can be two or more managed disks from the same managed disk group. The same circular algorithm is used across the striped set. However, a single managed disk can be specified more than once in the list. For example, if you enter **-m 0:1:2:1**, the extents are from the following **managed** disks: 0, 1, 2, 1, 0, 1, 2, and so forth. All MDisks that are specified in the **-mdisk** parameter must be in **managed** mode.

image This policy allows image mode virtual disks to be created when a managed disk already has data on it, perhaps from a previrtualized subsystem. When an image mode virtual disk is created, it directly corresponds to the (previously unmanaged) managed disk that it was created from; therefore, virtual disk logical block address (LBA) *x* equals managed disk LBA *x*. You can use this command to bring a nonvirtualized disk under the control of the cluster. After it is under the control of the cluster, you can migrate the virtual disk from the single managed disk. When it is migrated, the virtual disk is no longer an image mode virtual disk.

You can add image mode VDisks to an already populated MDisk group with other types of VDisks, such as a striped or sequential.

Note: An image mode copy must be at least as large as the VDisk that it is being added to, but any capacity beyond the size of the VDisk is not accessible.

The command returns the ID of the newly created VDisk copy.

Table 12 provides the relationship of the *rate* value to the data copied per second.

Table 12. Relationship between the rate value and the data copied per second

User-specified <i>rate</i> attribute value	Data copied/sec
1 - 10	128 KB
11 - 20	256 KB
21 - 30	512 KB
31 - 40	1 MB
41 - 50	2 MB
51 - 60	4 MB
61 - 70	8 MB
71 - 80	16 MB
81 - 90	32 MB
91 - 100	64 MB

An invocation example

```
| svctask addvdiskcopy -mdiskgrp 0 -easytier off vdisk8
```

The resulting output

```
Vdisk [8] copy [1] successfully created
```

An invocation example

```
| svctask addvdiskcopy -mdiskgrp 0 -vtype image -mdisk 13 -tier generic_ssd -easytier off vdisk9
```

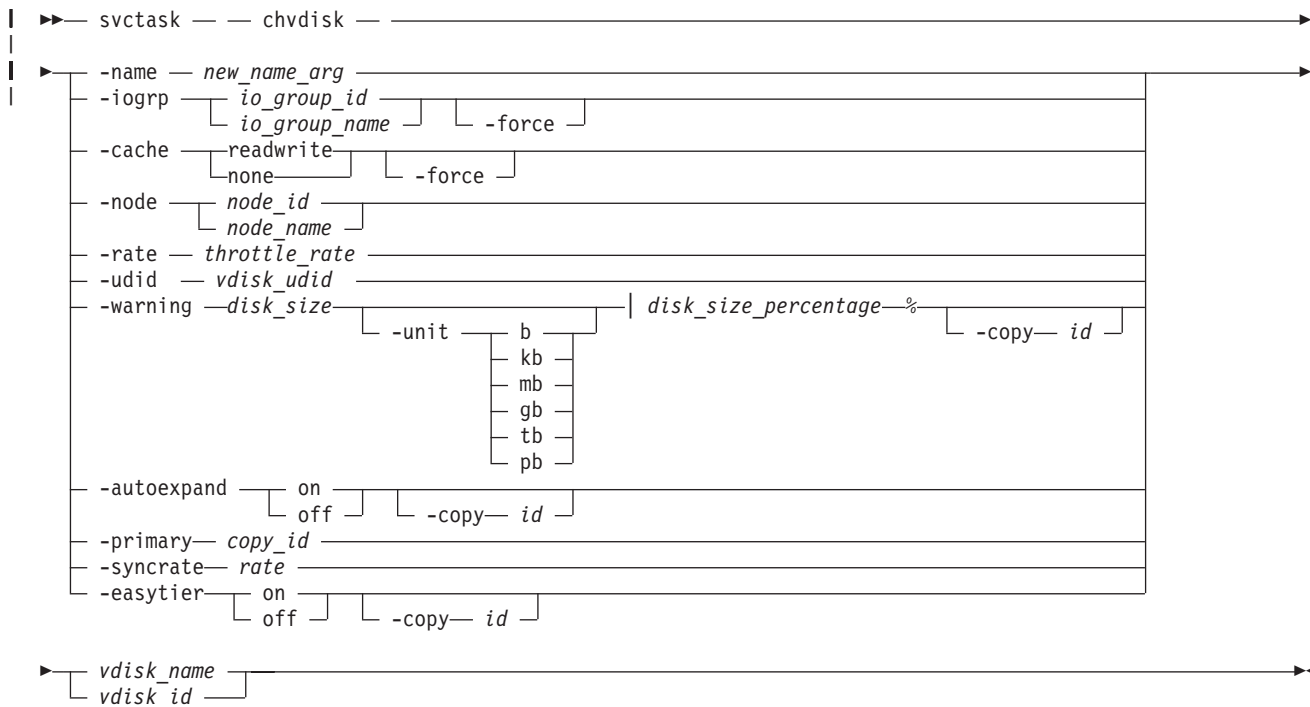
The resulting output

| Vdisk [9] copy [1] successfully created

chvdisk

| The **chvdisk** command modifies the properties of a VDisk (volume), such as the disk name, I/O group, I/O governing rate, or unit number.

Syntax



Parameters

-name *new_name_arg*

(Optional) Specifies a new name to assign to the virtual disk. You cannot use this parameter with the **-iogrp**, **-rate**, **-node**, or **-udid** parameters. This parameter is required if you do not use the **-iogrp**, **-rate**, or **-udid** parameter.

-iogrp *io_group_id* | *io_group_name*

(Optional) Specifies a new I/O group to move the virtual disk to, by I/O group ID or I/O group name. You can use the **-node** parameter with the **-iogrp** parameter to specify a preferred node for the specified VDisk (volume).

Notes:

1. If the VDisk (volume) has a mapping to any hosts, it is not possible to move the VDisk (volume) to an I/O group, unless all of those hosts are associated with the new I/O group.
2. This parameter can fail if there is not enough space to allocate bitmaps for a mirrored VDisk (volume) in the target I/O group.
3. This parameter can fail if any copy is not synchronized. The **-force** parameter can be used to force the move, but this resynchronizes the VDisk (volume).

-cache *readwrite* | *none*

(Optional) Specifies the caching options for the VDisk (volume). Valid entries are **readwrite**, to enable the cache for the VDisk (volume), or **none**, to disable the cache mode for the VDisk (volume).

-force

| (Optional) The **force** parameter can only be used for changing the I/O group of a VDisk (volume) or
| the caching mode. Use the **force** parameter with the **iogrp** parameter to force the VDisk (volume) to
| be removed from an I/O group. Use the **force** parameter with the **cache** parameter to specify that
| you want the system to change the cache mode of the VDisk (volume) even if the I/O group is
| offline. This option overrides the cache flush mechanism.

Attention:

- | 1. If the **force** parameter is used for changing the caching mode or I/O group of a VDisk (volume),
| the contents of the cache are discarded and the VDisk (volume) might be corrupted by the loss of
| the cached data. This could occur if the cluster is able to destage all write data from the cache or
| not. The **force** parameter should be used with caution.
- | 2. If the **force** parameter is used to move a VDisk (volume) that has out-of-sync copies, a full
| resynchronization is required.

-node *node_id* | *node_name*

| (Optional) Specifies a preferred node for the specified VDisk (volume). This parameter can only be
| specified when using the **iogrp** parameter to move the VDisk (volume) to another I/O group. This
| parameter is mutually exclusive to all other parameters except for **iogrp** and **-force**.

-rate *throttle_rate* [-**unitmb**]

| (Optional) Specifies the I/O governing rate for the VDisk (volume), which caps the amount of I/O
| that is accepted. The default *throttle_rate* units are I/Os. To change the *throttle_rate* units to megabytes
| per second (MBps), specify the **-unitmb** parameter. The governing rate for a virtual disk can be
| specified by I/Os or by MBps, but not both. However, you can set the rate to I/Os for some virtual
| disks and to MBps for others.

You cannot use this parameter with the **-name**, **-iogrp**, **-node**, or **-udid** parameters.

-udid *vdisk_udid*

(Optional) Specifies the unit number (**udid**) for the disk. The *vdisk_udid* is an identifier that is
required to support OpenVMS hosts; no other systems use this parameter. Valid options are a decimal
number from 0 to 32 767 or a hexadecimal number from 0 to 0x7FFF. A hexadecimal number must be
preceded by 0x (for example, 0x1234). If you do not use the **-udid** parameter, the default **udid** is 0.

You cannot use this parameter with the **-name**, **-iogrp**, **-node**, or **-rate** parameters.

-warning *disk_size* | *disk_size_percentage*%

(Optional) Generates a warning when the used disk capacity on the space-efficient copy first exceeds
the specified threshold. You can specify a *disk_size* integer, which defaults to MBs unless the **-unit**
parameter is specified; or you can specify a *disk_size*%, which is a percentage of the virtual disk size.
To disable warnings, specify 0 or 0%.

-unit **b** | **kb** | **mb** | **gb** | **tb** | **pb**

(Optional) Specifies the data units to use for the **-warning** *disk_size* parameter.

-autoexpand **on** | **off**

| (Optional) Specifies whether space-efficient VDisk (volume) copies automatically expand their real
| capacities by allocating new extents from their managed disk group. To use this parameter, the VDisk
| (volume) must be space-efficient.

-copy *id*

| (Optional) Specifies the copy to apply the changes to. You must specify this parameter with the
| **-autoexpand** or **-warning** parameter. The **-copy** parameter is required if the specified VDisk (volume)
| is mirrored and only one VDisk (volume) copy is space-efficient. If both copies are space-efficient and
| the **-copy** parameter is not specified, the specified **-autoexpand** or **-warning** parameter is set on both
| copies.

-primary *copy_id*

(Optional) Specifies the primary copy. Changing the primary copy only takes effect when the new

primary copy is online and synchronized. If the new primary is online and synchronized when the command is issued, the change takes effect immediately.

- | **-syncrate** *rate*
- | (Optional) Specifies the copy synchronization rate, as a percentage of the peak synchronization rate. A value of zero (0) prevents synchronization.
- | **-easytier on | off**
- | (Optional) Enables or disables the IBM System Storage Easy Tier function.
- | *vdisk_name | vdisk_id*
- | (Required) Specifies the VDisk (volume) to modify, either by ID or by name.

Description

- | The **chvdisk** command modifies a single property of a VDisk (volume). To change the VDisk (volume) name and modify the I/O group, for example, you must issue the command twice.

- | **Note:** If the VDisk (volume) is offline, use one of the **recovervdisk** commands to recover the VDisk (volume) and bring it back online.

- | You can specify a new name or label. You can use the new name subsequently to refer to the virtual disk.
- | To specify a preferred node for the VDisk (volume), use the **-node** *node_id | node_name* parameter.

You can change the I/O group with which this virtual disk is associated. However, to change the I/O group, you must first flush the cache within the nodes in the current I/O group to ensure that all data is written to disk. Ensure that you suspend I/O operations at the host level before you perform this operation.

Attention:

- | 1. Do not move a VDisk (volume) to an offline I/O group under any circumstance. To avoid any data loss, you must ensure that the I/O group is online before you move the VDIsks (volumes).
- | 2. Do not move an offline VDisk (volume) to the recovery I/O group. Beginning with SAN Volume Controller version 4.3.1, use of the recovery I/O group is not required. Instead, use one of the **recovervdisk** commands to recover the VDisk (volume) and bring it back online.

You can set a limit on the amount of I/O transactions that is accepted for this virtual disk. It is set in terms of I/Os per second or MBs per second. By default, no I/O governing rate is set when a virtual disk is created.

Attention: All capacities, including changes, must be in multiples of 512 bytes. An error occurs if you specify a capacity that is not a multiple of 512, which can only happen when byte units (**-b**) are used. The default capacity is in MB.

When the virtual disk is created, there is no throttling applied to it. Using the **-rate** parameter can change this. To change the virtual disk back to an unthrottled state, specify 0 (zero) with the **-rate** parameter.

- | You can migrate a VDisk (volume) to a new I/O group to manually balance the workload across the nodes in the cluster. You might end up with a pair of nodes that are overworked and another pair that are underworked. Use the following procedure to migrate a single VDisk (volume) to a new I/O group.
- | Repeat for other VDIsks (volumes) as required.

- | **Attention:** This is a disruptive procedure. Access to the VDisk (volume) is lost while you follow this procedure.

- | Ensure that when you migrate a VDisk (volume) to a new I/O group, you quiesce all I/O operations for the VDisk (volume). Determine the hosts that are using this VDisk (volume). Stop and delete any

| FlashCopy mappings or Metro or Global Mirror relationships that use this VDisk (volume). To check if
 | the VDisk (volume) is part of a relationship or mapping, issue the **svcinfo lsvdisk** *vdiskname* | *id*
 | command, where *vdiskname* | *id* is the name or ID of the VDisk (volume).

| Look for the **FC_id** and **RC_id** fields. If these are not blank, the VDisk (volume) is part of a mapping or
 | relationship. See the FlashCopy commands or Metro Mirror and Global Mirror commands for details on
 | how to stop or delete the mapping or relationship. Issue the following command to migrate the VDisk
 | (volume):

```
svctask chvdisk -iogrp newiogrpname|id vdiskname|id
```

Follow the procedure to discover the new vpaths and to check that each vpath is presenting the correct number of paths. Refer to Multipath Subsystem Device Driver (SDD) documentation for details on how to dynamically reconfigure SDD for the given host operating system.

| **Note:** The command fails if you attempt to change the primary copy of a mirrored VDisk (volume) while the **repairvdiskcopy -resync** command is running.

An invocation example

```
svctask chvdisk -rate 2040 -unitmb 6
```

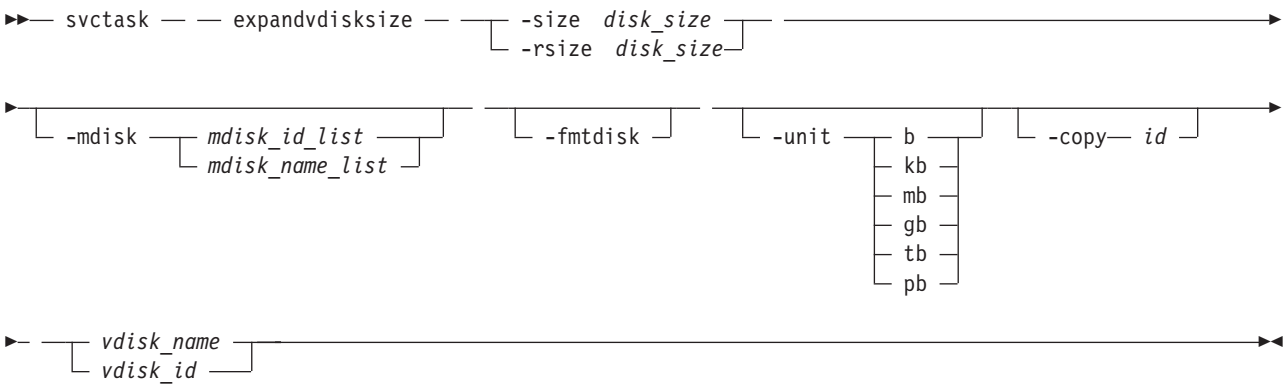
The resulting output

No feedback

expandvdisksize

The **expandvdisksize** command expands the size of a VDisk by a given capacity.

Syntax



Parameters

-size *disk_size*
 (Optional) Specifies the capacity by which the virtual disk is expanded. Disk size is used with the value of the unit. All capacities, including changes must be in multiples of 512 bytes. An error occurs if you specify a capacity that is not a multiple of 512, which can only occur when byte units (**-unit b**) are used. However, an entire extent is reserved even if it is only partially used. The default *disk_size* unit is megabytes (MB). You cannot specify the **-size** parameter with the **-rsize** parameter. You must specify either **-size** or **-rsize**. If the VDisk is space-efficient, MDisks cannot be specified.

-rsize *disk_size*
 (Optional) Specifies the capacity by which to increase the real size of a space-efficient VDisk. Specify

the *disk_size* value using an integer. Specify the unit for a *disk_size* integer using the **-unit** parameter; the default unit is megabytes (MB). The **-rsize** value can be greater than, equal to, or less than the size of the VDisk. You cannot specify the **-rsize** parameter with the **-size** parameter. You must specify either **-size** or **-rsize**.

-copy *id*

(Optional) Specifies the copy to change the real capacity for. You must also specify the **-rsize** parameter; you can only modify the real capacity of a VDisk copy. The **-copy** parameter is required if the specified VDisk is mirrored and only one copy is space-efficient. If the VDisk is mirrored, both copies are space-efficient and **-copy** is not specified, both copies are modified by the same amount.

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Optional) Specifies the list of one or more MDisks to be used as the stripe set. The extents that expand the VDisk come from the specified list of MDisks. All MDisks in the list must be part of the same MDisk group. The **-mdisk** parameter cannot be used if the specified VDisk is mirrored.

-fmtdisk

(Optional) Specifies that the VDisk be formatted before use. This parameter formats the new extents that have been added to the VDisk as a result of the **expandvdisksize** command. The **expandvdisksize** command completes asynchronously if you use this parameter.

-unit *b* | *kb* | *mb* | *gb* | *tb* | *pb*

(Optional) Specifies the *disk_size* unit for the **-size** or **-rsize** parameter. The default value is megabytes (MB).

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk to modify, either by ID or by name.

Description

The **expandvdisksize** command can be used to expand the physical capacity that is allocated to a particular VDisk by the specified amount. The command can also be used to expand the virtual capacity of a space-efficient VDisk without altering the physical capacity that is assigned to the VDisk. To change the capacity of a non-space-efficient VDisk, or the virtual capacity of a space-efficient VDisk, use the **-size** parameter. To change the real capacity of a space-efficient VDisk, use the **-rsize** parameter.

The default capacity units are MB.

When a VDisk is expanded, the virtualization policy can change. Its mode becomes striped even if it was previously sequential. See the **svctask mkvdisk** command for details of the virtualization policies.

To run the **expandvdisksize** command on a mirrored VDisk, all copies of the VDisk must be synchronized. The command formats all copies of a mirrored VDisk automatically.

An invocation example

To increase the capacity of VDisk1 by 2048 bytes by using extents from two MDisks and to format the new part of the VDisk, enter:

```
svctask expandvdisksize -size 2048 -unit b -mdisk mdisk0:mdisk1 -fmtdisk vdisk1
```

The resulting output

No feedback

An invocation example

To increase the capacity of VDisk1 by 100 MB using extents from two MDisks, and to format the new part of the VDisk, enter:

```
svctask expandvdisksize -size 100 -unit mb -mdisk mdisk0:mdisk1 -fmtdisk vdisk1
```

The resulting output

No feedback

An invocation example

To increase the real capacity of space-efficient VDisk **vdisk2** by 100 MB without changing the virtual capacity, and to spread the extents across all MDisk in the MDisk group, enter:

```
svctask expandvdisksize -rsize 100 -unit mb vdisk2
```

The resulting output

No feedback

An invocation example

To increase the real capacity of space-efficient VDisk copy id 1 of mirrored VDisk **vdisk3** by 100 MB, enter:

```
svctask expandvdisksize -rsize 100 -unit mb -copy 1 vdisk3
```

The resulting output

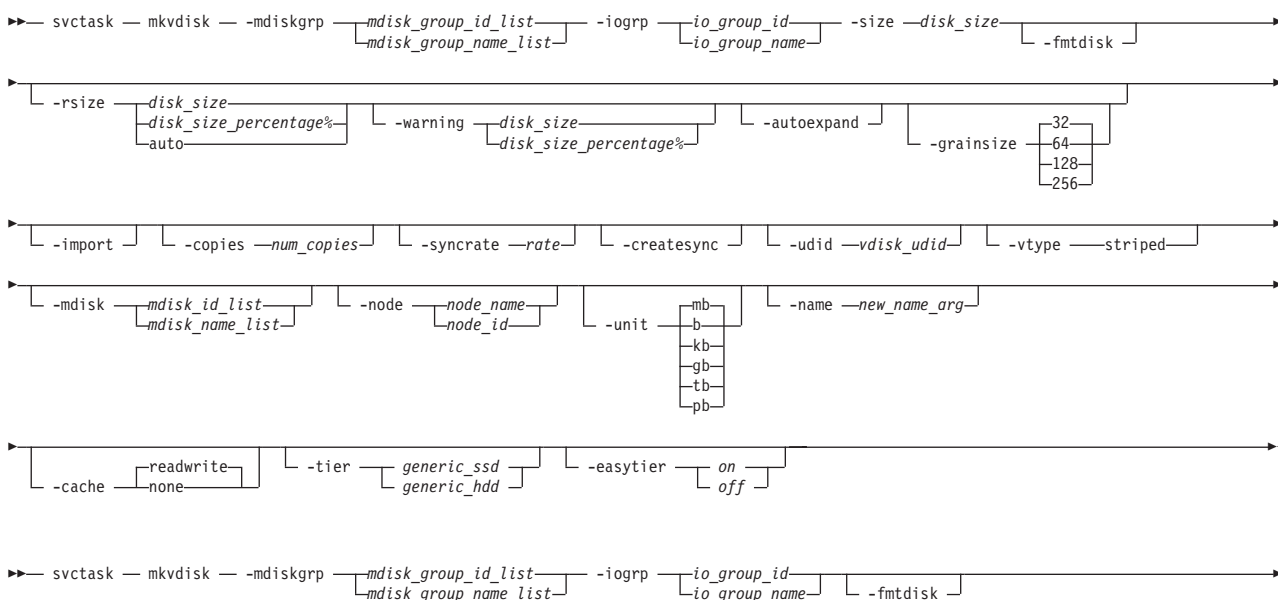
No feedback

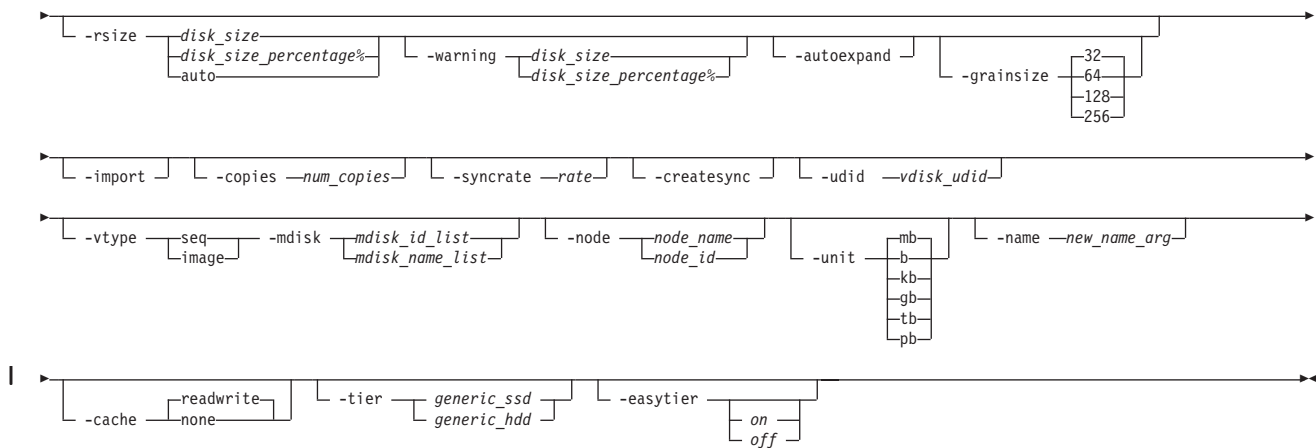
mkvdisk

- 1 The **mkvdisk** command creates sequential, striped, or image mode virtual disk (volume) objects. When they are mapped to a host object, these objects are seen as disk drives with which the host can perform I/O operations.

Note: The first syntax diagram depicts the creation of a **striped** mode virtual disk. The second syntax diagram depicts the creation of a **sequential** or **image** mode virtual disk.

Syntax





Parameters

-mdiskgrp *mdisk_group_id_list* | *mdisk_group_name_list*

(Required) Specifies one or more managed disk groups (storage pools) to use when you are creating this virtual disk. If you are creating multiple copies, you must specify one managed disk group per copy. The primary copy is allocated from the first managed disk group in the list.

-iogrp *io_group_id* | *io_group_name*

(Required) Specifies the I/O group (node pair) with which to associate this virtual disk.

-udid *vdisk_udid*

(Optional) Specifies the unit number (**udid**) for the disk. The **udid** is an identifier that is required to support OpenVMS hosts; no other systems use this parameter. Valid options are a decimal number 0 - 32 767, or a hexadecimal number 0 - 0x7FFF. A hexadecimal number must be preceded by **0x** (for example, **0x1234**).

-size *disk_size*

(Required for sequential [**seq**] or **striped** VDisk creation) (Optional for **image** VDisk creation) Specifies the capacity of the virtual disk, which is used with the value of the unit. All capacities, including changes, must be in multiples of 512 bytes. An error occurs if you specify a capacity that is not a multiple of 512, which can only happen when byte units (**-b**) are used. However, an entire extent is reserved even if it is only partially used. The default capacity is in MB. You can specify a capacity of 0. Specify the size in bytes in multiples of logical block address (LBA) sizes.

Note: If you do not specify the **-size** parameter when you create an image mode disk, the entire MDisk capacity is used.

-rsize *disk_size* | *disk_size_percentage%* | **auto**

(Optional) Defines how much physical space is initially allocated to the space-efficient VDisk (thin-provisioned volume). This parameter makes the VDisk space-efficient; otherwise, the VDisk is fully allocated. Specify the *disk_size* | *disk_size_percentage* value using an integer, or an integer immediately followed by the percent character (%). Specify the units for a *disk_size* integer using the **-unit** parameter; the default is MB. The **-rsize** value can be greater than, equal to, or less than the size of the VDisk. The **auto** option creates a VDisk copy that uses the entire size of the MDisk; if you specify the **-rsize auto** option, you must also specify the **-vtype image** option.

-warning *disk_size* | *disk_size_percentage%*

(Optional) Requires that the **-rsize** parameter also be specified. Specifies a threshold at which a warning error log is generated for VDisk copies. A warning is generated when the used disk capacity on the space-efficient copy first exceeds the specified threshold. You can specify a *disk_size* integer, which defaults to MBs unless the **-unit** parameter is specified; or you can specify a *disk_size%*, which

is a percentage of the virtual disk size. If **-autoexpand** is enabled, the default value for **-warning** is 80% of the virtual disk capacity. If **-autoexpand** is not enabled, the default value for warning is 80% of the real capacity. To disable warnings, specify **0**.

-autoexpand

(Optional) Specifies that space-efficient copies automatically expand their real capacities by allocating new extents from their managed disk group. Requires that the **-rsize** parameter also be specified. If the **-autoexpand** parameter is specified, the **-rsize** parameter specifies a capacity that is reserved by the copy. This protects the copy from going offline when its managed disk group runs out of space by having the managed disk group to consume this reserved space first.

The parameter has no immediate effect on **image** mode copies. However, if the image mode copy is subsequently migrated to managed mode, the copy is then automatically expanded.

-grainsize 32 | 64 | 128 | 256

(Optional) Sets the grain size (KB) for a space-efficient VDisk. This parameter also requires that the **-rsize** parameter be specified. The default is 32 KB. If you are using the space-efficient VDisk in a FlashCopy map, use the same grain size as the map grain size for best performance. If you are using the space-efficient VDisk directly with a host system, use a small grain size.

-import

(Optional) Imports a space-efficient VDisk from the Mdisk. This parameter also requires that the **-rsize** parameter be specified.

-copies *num_copies*

(Optional) Specifies the number of copies to create. The *num_copies* value can be 1 or 2. Setting the value to 2 creates a mirrored VDisk. The default value is 1.

-syncrate *rate*

(Optional) Specifies the copy synchronization rate. A value of zero (**0**) prevents synchronization. The default value is **50**. For the supported **-syncrate** values and their corresponding rates, see Table 14 on page 165.

-createsync

(Optional) Creates copies in sync. Use this parameter if you have already formatted the MDisks, or when read stability to unwritten areas of the VDisk is not required.

-fmtdisk

(Optional) Specifies that the virtual disk be formatted before it can be used. The **-fmtdisk** parameter formats (sets to all zeros) the extents that make up this VDisk after it is created. If this parameter is used, the command completes asynchronously; you can query the status using the **svcinfolsvdiskprogress** command.

The **-fmtdisk** parameter is not required when creating space-efficient virtual disks. Space-efficient VDIsks return zeros for extents that have not been written to.

The **-fmtdisk** parameter synchronizes mirrored copies by default.

Note: You cannot specify this parameter with the **-vtype image** parameter.

-vtype seq | striped | image

(Optional) Specifies the virtualization type. When creating sequential or image mode VDIsks, you must also specify the **-mdisk** parameter. The default virtualization type is striped.

-node *node_id* | *node_name*

(Optional) Specifies the preferred node ID or the name for I/O operations to this virtual disk. You can use the **-node** parameter to specify the preferred access node.

Note: This parameter is required for the subsystem device driver (SDD). The cluster chooses a default if you do not supply this parameter.

-unit *b | kb | mb | gb | tb | pb*

(Optional) Specifies the data units to use in conjunction with the capacity that is specified by the **-size** and **-rsize** parameters.

-mdisk *mdisk_id_list | mdisk_name_list*

(Optional) Specifies one or more managed disks. For sequential and image mode VDisks, the number of MDisks must match the number of copies. For sequential mode VDisks, each MDisk must belong to the specified MDisk group. For striped VDisks, you cannot specify the **-mdisk** parameter if the **-copies** value is greater than 1. When creating a single copy striped VDisk, you can specify a list of MDisks to stripe across.

-name *new_name_arg*

(Optional) Specifies a name to assign to the new virtual disk.

-cache *readwrite | none*

(Optional) Specifies the caching options for the VDisk. Valid entries are **readwrite** or **none**. The default is **readwrite**. If you do not specify the **-cache** parameter, the default value (**readwrite**) is used.

-tier

(Optional) Specifies the MDisk tier when an image mode copy is added.

Note: This will apply to both copies if you are creating mirrored VDisk with two image mode copies using this command.

-easytier

(Optional) Determines if the IBM System Storage Easy Tier function is allowed to move extents for this VDisk. If a volume copy is striped and not being migrated the following table applies:

Table 13. Storage pool Easy Tier settings

Storage pool Easy Tier setting	Number of tiers in the storage pool	Volume copy Easy Tier setting	Volume copy Easy Tier status
Off	One	Off	inactive (see note 2 on page 164)
Off	One	On	inactive (see note 2 on page 164)
Off	Two	Off	inactive (see note 2 on page 164)
Off	Two	On	inactive (see note 2 on page 164)
Auto (see note 5 on page 164)	One	Off	inactive (see note 2 on page 164)
Auto (see note 5 on page 164)	One	On	inactive (see note 2 on page 164)
Auto (see note 5 on page 164)	Two	Off	measured (see note 3 on page 164)
Auto (see note 5 on page 164)	Two	On	active (see note 1 on page 164)
On (see note 5 on page 164)	One	Off	measured (see note 3 on page 164)
On (see note 5 on page 164)	One	On	measured (see note 3 on page 164)
On (see note 5 on page 164)	Two	Off	measured (see note 3 on page 164)
On (see note 5 on page 164)	Two	On	active (see note 1 on page 164)

Table 13. Storage pool Easy Tier settings (continued)

Storage pool Easy Tier setting	Number of tiers in the storage pool	Volume copy Easy Tier setting	Volume copy Easy Tier status
Note:			
1. If the volume copy is in image or sequential mode or is being migrated then the volume copy Easy Tier status will be <i>measured</i> instead of <i>active</i> .			
2. When the volume copy status is <i>inactive</i> , no Easy Tier functions are enabled for that volume copy.			
3. When the volume copy status is <i>measured</i> , the Easy Tier function collects usage statistics for the volume but automatic data placement is not active.			
4. When the volume copy status is <i>active</i> , the Easy Tier function operates in automatic data placement mode for that volume.			
5. The default Easy Tier setting for a storage pool is <i>auto</i> , and the default Easy Tier setting for a volume copy is <i>on</i> . This means that Easy Tier functions will be disabled for storage pools with a single tier, and that automatic data placement mode will be enabled for all striped volume copies in a storage pool with two tiers.			

Description

This command creates a new virtual disk object. You can use the command to create a variety of types of virtual disk objects, making it one of the most complex commands.

You must decide which managed disk group or groups provide the storage for the VDisk. Use the **svcinfolsmdiskgrp** command to list the available managed disk groups and the amount of free storage in each group. If you are creating a VDisk with more than one copy, each MDisk group that you specify must have enough space for the size of the VDisk.

Important: The extent size for the MDisk group can limit VDisk size. Consider the maximum VDisk size you want to use when creating MDisk groups. Refer to the information on creating MDisk groups for a comparison of the maximum VDisk capacity for each extent size. The maximum is different for space-efficient Vdisks (thin-provisioned volumes).

Choose an I/O group for the VDisk. This determines which nodes in the cluster process the I/O requests from the host systems. If you have more than one I/O group, ensure that you distribute the VDisks between the I/O groups so that the I/O workload is shared evenly between all nodes. Use the **svcinfolsiogrp** command to show the I/O groups and the number of virtual disks that are assigned to each I/O group.

Note: It is normal for clusters with more than one I/O group to have MDisk groups that have VDisks in different I/O groups. FlashCopy processing can make copies of VDisks whether the source and target VDisks are in the same I/O group. If, however, you plan to use intracluster Metro or Global Mirror operations, ensure that both the master and auxiliary VDisk are in the same I/O group.

Specify the virtualization type using the **-vtype** parameter; the supported types are **sequential** (**seq**), **striped**, and **image**.

sequential (seq)

This virtualization type creates the virtual disk using sequential extents from the specified MDisk (or MDisks, if creating multiple copies). The command fails if there are not enough sequential extents on the specified MDisk.

striped

This is the default virtualization type. If the **-vtype** parameter is not specified, **striped** is the default; all managed disks in the managed disk group are used to create the virtual disk. The striping is at an extent level; one extent from each managed disk in the group is used. For example, a managed disk group with 10 managed disks uses one extent from each managed disk, then it uses the 11th extent from the first managed disk, and so on.

If the **-mdisk** parameter is also specified, you can supply a list of managed disks to use as the stripe set. This can be two or more managed disks from the same managed disk group. The same circular algorithm is used across the striped set. However, a single managed disk can be specified more than once in the list. For example, if you enter **-mdisk 0:1:2:1**, the extents are from the following managed disks: 0, 1, 2, 1, 0, 1, 2, and so forth. All MDisks that are specified in the **-mdisk** parameter must be in the managed mode.

A capacity of 0 is allowed.

image This virtualization type allows image mode virtual disks to be created when a managed disk already has data on it, perhaps from a previrtualized subsystem. When an image mode virtual disk is created, it directly corresponds to the (previously unmanaged) managed disk that it was created from. Therefore, with the exception of space-efficient image mode VDIs, virtual disk logical block address (LBA) *x* equals managed disk LBA *x*. You can use this command to bring a nonvirtualized disk under the control of the cluster. After it is under the control of the cluster, you can migrate the virtual disk from the single managed disk. When it is migrated, the virtual disk is no longer an image mode virtual disk.

You can add image mode VDIs to an already populated MDisk group with other types of VDIs, such as a striped or sequential.

Note: An image mode VDisk must be 512 bytes or greater. At least one extent is allocated to an image mode VDisk.

You must use the **-mdisk** parameter to specify an MDisk that has a mode of unmanaged. The **-fmtmdisk** parameter cannot be used to create an image mode VDisk.

Note: If you create a mirrored VDisk from two image mode MDisks without specifying a **-capacity** value, the capacity of the resulting VDisk is the smaller of the two MDisks, and the remaining space on the larger MDisk is not accessible.

The command returns the IDs of the newly created VDisk.

Attention:

1. Do not create a VDisk in an offline I/O group. You must ensure that the I/O group is online before you create a VDisk to avoid any data loss. This applies in particular to recreating VDIs that are assigned the same object ID.
2. To create an image mode disk, you must already have a quorum disk in the cluster because an image mode disk cannot be used to hold quorum data. Refer to information on quorum disk creation for more details.
3. The command fails if either limit of 2048 VDIs per I/O Group or 8192 VDisk copies per cluster is reached.

Table 14 provides the relationship of the *rate* value to the data copied per second.

Table 14. Relationship between the rate value and the data copied per second

User-specified <i>rate</i> attribute value	Data copied/sec
1 - 10	128 KB
11 - 20	256 KB
21 - 30	512 KB
31 - 40	1 MB
41 - 50	2 MB
51 - 60	4 MB
61 - 70	8 MB

Table 14. Relationship between the rate value and the data copied per second (continued)

User-specified <i>rate</i> attribute value	Data copied/sec
71 - 80	16 MB
81 - 90	32 MB
91 - 100	64 MB

An invocation example

```
svctask mkvdisk -mdiskgrp Group0 -size 0
-iogrp 0 -vtype striped -mdisk mdisk1 -node 1
```

The resulting output

Virtual Disk, id [1], successfully created

An invocation example for creating an image mode VDisk

```
svctask mkvdisk -mdiskgrp Group0
-iogrp 0 -vtype image -mdisk mdisk2 -node 1
```

The resulting output

Virtual Disk, id [2], successfully created

An invocation example for creating a new VDisk

```
svctask mkvdisk -mdiskgrp Group0 -size 0 -unit kb
-iogrp 0 -vtype striped -mdisk mdisk1 -node 1 -udid 1234 -easytier off
```

The resulting output

Virtual Disk id [2], successfully created

An invocation example for creating a space-efficient VDisk

```
svctask mkvdisk -mdiskgrp Group0 -iogrp 0 -vtype striped
-size 10 -unit gb -rsize 20% -autoexpand -grainsize 32
```

The resulting output

Virtual Disk id [1], successfully created

An invocation example for creating a mirrored image-mode VDisk

```
svctask mkvdisk -mdiskgrp Group0:Group0 -mdisk mdisk2:mdisk3
-iogrp 0 -vtype image -copies 2
```

The resulting output

Virtual Disk id [1], successfully created

An invocation example for creating a mirrored VDisk

```
svctask mkvdisk -iogrp 0 -mdiskgrp 0:1 -size 500 -copies 2
```


The resulting output

Virtual Disk id [5], successfully created

mkvdiskhostmap

The **mkvdiskhostmap** command creates a new mapping between a virtual disk and a host, which makes the virtual disk accessible for I/O operations to the specified host.

Syntax



Parameters

-force

(Optional) Allows multiple VDisk-to-host assignments, which are not normally allowed.

-host *host_id* | *host_name*

(Required) Specifies the host to map the virtual disk to, either by ID or by name.

-scsi *scsi_num_arg*

(Optional) Specifies the SCSI LUN ID to assign to this virtual disk on the given host. The *scsi_num_arg* parameter contains the SCSI LUN ID that is assigned to the VDisk on the given host. You must check your host system for the next available SCSI LUN ID on the given HBA. If you do not specify the **-scsi** parameter, the next available SCSI LUN ID is provided to the host.

vdisk_name | *vdisk_id*

(Required) Specifies the name of the virtual disk that you want to map to the host, either by ID or by name.

Description

This command creates a new mapping between the virtual disk and the specified host. The virtual disk is presented to the host as if the disk is directly attached to the host. It is only after this command is processed, that the host can perform I/O transactions to the virtual disk.

Optionally, you can assign a SCSI LUN ID to the mapping. When the HBA in the host scans for devices that are attached to it, it discovers all virtual disks that are mapped to its fibre-channel ports. When the devices are found, each one is allocated an identifier (SCSI LUN ID). For example, the first disk found is usually SCSI LUN 1, and so on. You can control the order in which the HBA discovers virtual disks by assigning the SCSI LUN ID, as required. If you do not specify a SCSI LUN ID, the cluster automatically assigns the next available SCSI LUN ID, if any mappings already exist with that host. When you issue the **mkvdiskhostmap** command, the assigned SCSI LUN ID number is returned.

Some HBA device drivers will stop when they find a gap in the SCSI LUN IDs. For example:

- Virtual Disk 1 is mapped to Host 1 with SCSI LUN ID 1
- Virtual Disk 2 is mapped to Host 1 with SCSI LUN ID 2
- Virtual Disk 3 is mapped to Host 1 with SCSI LUN ID 4

When the device driver scans the HBA, it must stop after identifying virtual disks 1 and 2, because no SCSI LUN is mapped with ID 3. For optimal performance, ensure that the SCSI LUN ID allocation is contiguous.

You can create multiple VDisk assignments. Normally, multiple VDisk-to-host assignments are not used because corruption is likely to occur if more than one host can access a disk. However, in certain multiple path environments, such as in the IBM SAN File System, a VDisk must be mapped to more than one host. To map to more than one host, you must use the **mkvdiskhostmap** command with the **force** parameter. For example:

```
svctask mkvdiskhostmap -host host1 -force 4
svctask mkvdiskhostmap -host host2 -force 4
```

These commands create two host-to-VDisk mappings for VDisk 4 that map to host1 and host2. Omitting the **force** parameter causes the mapping to fail if that VDisk is already mapped to a host.

The command also fails if the host object (to which this mapping is being made) is not associated with the I/O group containing the VDisk.

An invocation example

```
svctask mkvdiskhostmap -host host1 -scsi 1 5
```

The resulting output

```
Virtual Disk to Host map, id [1], successfully created
```

recovervdisk

The **recovervdisk** command acknowledges VDisk data loss and brings the VDisk back online.

Syntax

```
| ▶▶—svctask— —recovervdisk— —copy— —copy_id— —vdisk_name— —vdisk_id—▶▶
```

Parameters

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk to recover.

| **-copy** *copy_id*

| (Optional) Specifies the ID of the copy to recover.

Description

The specified VDisk, and all copies if mirrored, are recovered and brought back online. If the VDisk is space-efficient or has space-efficient copies, this command triggers the space-efficient repair process. If the VDisk is mirrored, the **recovervdisk** command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored using the **svcinfolsvdisksyncprogress** command. The VDisk remains online during the resynchronization process.

The **recovervdisk** command also starts the repair of any space-efficient copies that have a `fast_write_state` of corrupt. The progress of the repair process can be monitored using the **svcinfolrepairsevdiskcopyprogress** command.

A VDisk that is still offline because it is being repaired following the **recovervdisk** command has a `fast_write_state` of repairing. The VDisk is brought online when the repair process is complete.

An invocation example (to recover VDisk 45)

```
| svctask recovervdisk vdisk45
```

An invocation example (to recover copy 0 of VDisk 45)

```
| svctask recovervdisk -copy 0 vdisk45
```

recovervdiskbycluster

The **recovervdiskbycluster** command acknowledges data loss for all VDIs in the cluster with a `fast_write_state` of corrupt and brings the VDIs back online.

Syntax

```
▶▶—svctask— —recovervdiskbycluster—
```

Parameters

There are no parameters.

Description

All VDIs in the cluster that have a `fast_write_state` of corrupt; and all copies, if mirrored, are recovered and brought back online. If any of the VDIs are space-efficient or have space-efficient copies, the **recovervdiskbycluster** command triggers the space-efficient repair process. If VDIs are mirrored, the command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored by using the **svcinfolsvdisksyncprogress** command. VDIs remain online during the resynchronization process.

If none of the VDIs in the cluster have a `fast_write_state` of corrupt, the **recovervdiskbycluster** command still starts the repair process for any corrupt copies of mirrored VDIs. The progress of the repair process can be monitored using the **svcinfolrepairsevdiskcopyprogress** command. If there are no corrupt VDIs or no repairs to copies are required, no error is returned.

VDIs that are still offline because they are being repaired following the **recovervdiskbycluster** command have a `fast_write_state` of repairing. VDIs are brought online when the repair process is complete.

An invocation example

```
svctask recovervdiskbycluster
```

The resulting output

No feedback

recovervdiskbyiogrp

The **recovervdiskbyiogrp** command acknowledges data loss for all VDIs in the specified I/O group with a `fast_write_state` of corrupt and brings the VDIs back online.

Syntax

```
▶▶—svctask— —recovervdiskbyiogrp— io_group_name | io_group_id
```

Parameters

io_group_name | *io_group_id*
(Required) Specifies the I/O group for virtual disk recovery.

Description

All VDIs in the specified I/O group that have a `fast_write_state` of corrupt; and all copies, if mirrored, are recovered and brought back online. If any of the VDIs are space_efficient or have space_efficient copies, the **recovervdiskbyiogrp** command triggers the space-efficient repair process. If VDIs are

mirrored, the command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored by using the **svcinfolsvdisksyncprogress** command. VDIsks remain online during the resynchronization process.

If none of the VDIsks in the specified I/O group have a `fast_write_state` of `corrupt`, the **recovervdiskbyiogrp** command still starts the repair process for any corrupt copies of mirrored VDIsks. The progress of the repair process can be monitored using the **svcinfolrepairsevdiskcopyprogress** command. If there are no corrupt VDIsks or no repairs to copies are required, no error is returned.

VDIsks that are still offline because they are being repaired following the **recovervdiskbyiogrp** command have a `fast_write_state` of `repairing`. VDIsks are brought online when the repair process is complete.

An invocation example

```
svctask recovervdiskbyiogrp iogrp2
```

The resulting output

No feedback

repairsevdiskcopy

The **repairsevdiskcopy** command repairs the metadata on a space-efficient virtual disk.

Syntax

```
svctask repairsevdiskcopy [-copy 0 | 1] [vdisk_name | vdisk_id]
```

Parameters

-copy 0 | 1

(Optional) Specifies the VDisk copy to repair.

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk to repair.

Description

The **repairsevdiskcopy** command repairs the metadata on a space-efficient VDisk. Run this command only when you are directed by the Directed Maintenance Procedures or by IBM support.

Running the command automatically detects corrupted metadata. The command holds the VDisk offline during the repair, but does not prevent the disk from being moved between I/O groups.

If a repair operation completes successfully and the volume was previously offline because of corrupted metadata, the command brings the volume back online. The only limit on the number of concurrent repair operations is the number of virtual disk copies in the configuration. Once started, a repair operation cannot be paused or canceled; the repair can only be ended by deleting the copy.

An invocation example

```
svctask repairsevdiskcopy vdisk8
```


The resulting output

No feedback

repairvdiskcopy

- | The **repairvdiskcopy** command detects and optionally, corrects any VDisk (volume) copies that are not identical.

Syntax



Parameters

-medium

(Optional) Converts sectors that contain different contents into virtual medium errors on the specified VDisk. This parameter cannot be used with the **-validate** and **-resync** parameters; you must enter one of the three parameters.

-resync

(Optional) Corrects sectors that contain different contents by copying contents from the primary VDisk copy to other copies on the specified VDisk. This parameter cannot be used with the **-medium** and **-validate** parameters; you must enter one of the three parameters.

-validate

(Optional) Reports the first difference found on synchronized online copies of the specified VDisk, on or after the specified **-startlba** value. This parameter cannot be used with the **-medium** and **-resync** parameters; you must enter one of the three parameters.

-startlba lba

(Optional) Specifies a starting Logical Block Address (LBA) on which to begin the command.

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk to repair. You must specify this parameter last on the command line.

Description

The **repairvdiskcopy** command detects and optionally, corrects any VDisk copies that are not identical. The results are logged to the SAN Volume Controller error log. The **-validate** parameter compares synchronized online copies of the specified VDisk. The **-medium** parameter changes any sectors that are not identical into virtual medium errors. The **-resync** parameter copies any sectors that are not identical to the other VDisk copies. You must specify only one of the three parameters.

Attention:

1. Before you run the **repairvdiskcopy** command, ensure that all VDisk copies are synchronized.
2. Only one **repairvdiskcopy** command can run on a VDisk at a time. You must wait for the **repairvdiskcopy** command to complete processing before running the command again.
3. Once you start the **repairvdiskcopy** command, you cannot use the command to stop processing.
4. The primary copy of a mirrored VDisk cannot be changed while the **repairvdiskcopy -resync** command is running.

Use the **-startlba** parameter to specify a starting Logical Block Address (LBA). Enter an LBA value from 0 - full disk size minus one. The parameter logs the first error found and then stops the command. By repeating this parameter, you can collect all of the instances where the VDisk copies are not identical.

During **repairvdiskcopy** command operation, the VDisk remains online. The I/O and synchronization operations are allowed while the command is in progress.

The rate for the **repairvdiskcopy** command is controlled by the synchronization rate of the VDisk that is being repaired. To suspend the repair process, set the synchronization rate of the VDisk to **0** using the **chvdisk** command.

An invocation example

```
svctask repairvdiskcopy -resync vdisk8
```

The resulting output

No feedback

rmvdisk

The **rmvdisk** command deletes a virtual disk (VDisk).

Syntax

```
▶▶ svctask — — rmvdisk — — [ -force ] [ vdisk_id | vdisk_name ] ▶▶
```

Parameters

-force

(Optional) Deletes the specified VDisk, even if mappings still exist between this virtual disk and one or more hosts. This parameter deletes any host-to-VDisk mappings and any FlashCopy mappings that exist for this VDisk. This parameter also deletes any Metro Mirror or Global Mirror relationships that exist for the specified VDisk. Data that is on the virtual disk is lost. Before you issue this command, ensure that the virtual disk and any data that resides on it are no longer required.

vdisk_id | vdisk_name

Specifies the name of the virtual disk to delete, either by ID or by name.

Description

This command deletes an existing managed mode virtual disk or an existing image mode virtual disk. The extents that made up this virtual disk are returned to the pool of free extents that are available on the managed disk group, if the VDisk is in managed mode.

Attention: Any data that was on the virtual disk is lost. Before you issue this command, ensure that the virtual disk (and any data that resides on it) is no longer required.

Deleting a managed mode virtual disk

When you use this command to delete a managed mode virtual disk, all the data on the virtual disk is deleted. The extents that make up the virtual disk are returned to the pool of free extents that are available in the managed disk group.

If host mappings exist for the virtual disk, or if any FlashCopy mappings would be affected, the deletion fails. You can use the **-force** parameter to force the deletion. If you use the **-force** parameter, mappings that have the virtual disk as source or target are deleted, other mappings in a cascade might be stopped, and then the virtual disk is deleted. The **-force** parameter also deletes any Metro Mirror or Global Mirror relationships that exist for the specified VDisk.

If the virtual disk is in the process of migrating to an image mode virtual disk (using the `svctask migratetoimage` command), the deletion fails unless you use the `-force` parameter. If you use the `-force` parameter, the migration is halted and then the virtual disk is deleted. Before you issue this command, ensure that the virtual disk (and any data that resides on it) is no longer required.

Deleting an image mode virtual disk

If the VDisk is mirrored and one or both copies is in image mode, you must first wait for all fast-write data to be moved to the controller logical unit. This ensures that the data on the controller is consistent with the data on the image mode virtual disk before the VDisk is deleted. This process can take several minutes to complete, and is indicated by the `fast_write_state` state of the virtual disk being `empty`. If the `-force` parameter is specified, the fast-write data is discarded and the virtual disk is deleted immediately; the data on the controller logical unit is left inconsistent and unusable. If the copies are not synchronized, you must use the `-force` parameter.

If you run the command while data is in the cache, SVC attempts to move the data out of the cache; this process can time out, however.

If there are any virtual medium errors on the virtual disk, the command fails. You can force the deletion by using the `-force` parameter; however, this can cause data integrity problems.

Note: A virtual medium error occurs when you copy data from one disk (the source) to another (the target). Reading the source indicates that there is a medium error. At that moment, you must have two identical copies of data and you must then simulate a medium error on the target disk. You can simulate a medium error on the target disk by creating a virtual medium error on the target disk.

If FlashCopy mappings or host mappings exist for the virtual disk, the deletion fails unless you use the `-force` parameter. If you use the `-force` parameter, mappings are deleted and the virtual disk is deleted. If there is any data that is not staged in the fast write cache for this virtual disk, the deletion of the virtual disk fails. When the `-force` parameter is specified, any data that is not staged in the fast write cache is deleted. Deleting an image mode virtual disk causes the managed disk that is associated with the virtual disk to be removed from the managed disk group. The mode of the managed disk is returned to “unmanaged.”

An invocation example

```
svctask rmvdisk -force vdisk5
```

The resulting output

No feedback

rmvdiskcopy

The `rmvdiskcopy` command removes a VDisk copy from a VDisk.

Syntax

```
svctask --rmvdiskcopy --copy --copy_id [ -force ] [ vdisk_name vdisk_id ]
```

Parameters

`-copy copy_id`
(Required) Specifies the ID of the copy to delete.

`-force`
(Optional) Forces the deletion of the last synchronized copy of a VDisk, which deletes the entire

VDisk. The parameter also forces the deletion of a nonmirrored VDisk, a copy that is migrating to image mode, or an image-mode copy that has virtual medium errors.

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk to delete the copy from. You must specify this parameter last on the command line.

Description

The **rmvdiskcopy** command deletes the specified copy from the specified VDisk. The command fails if all other copies of the VDisk are not synchronized; in this case, you must specify the **-force** parameter, delete the VDisk, or wait until the copies are synchronized.

An invocation example

```
svctask rmvdiskcopy -copy 1 vdisk8
```

The resulting output

No feedback

rmvdiskhostmap

The **rmvdiskhostmap** command deletes an existing virtual disk-to-host mapping; the virtual disk is no longer accessible for I/O transactions on the given host.

Syntax

```
svctask -- rmvdiskhostmap -- -host [ host_id | host_name ] [ vdisk_id | vdisk_name ]
```

Parameters

-host *host_id* | *host_name*

(Required) Specifies the host that you want to remove from the map with the virtual disk, either by ID or by name.

vdisk_id | *vdisk_name*

(Required) Specifies the name of the virtual disk that you want to remove from the host mapping, either by ID or by name.

Description

This command deletes an existing mapping between the specified virtual disk and the host. This effectively stops the virtual disk from being available for I/O transactions on the given host.

This command also deletes a SCSI or persistent reservation that a host has on a VDisk. Once the reservation is removed, a new host is allowed to access the VDisk in the future because the original host no longer has access.

Use caution when you process this command because to the host, it seems as if the virtual disk has been deleted or is offline.

An invocation example

```
svctask rmvdiskhostmap -host host1 vdisk8
```

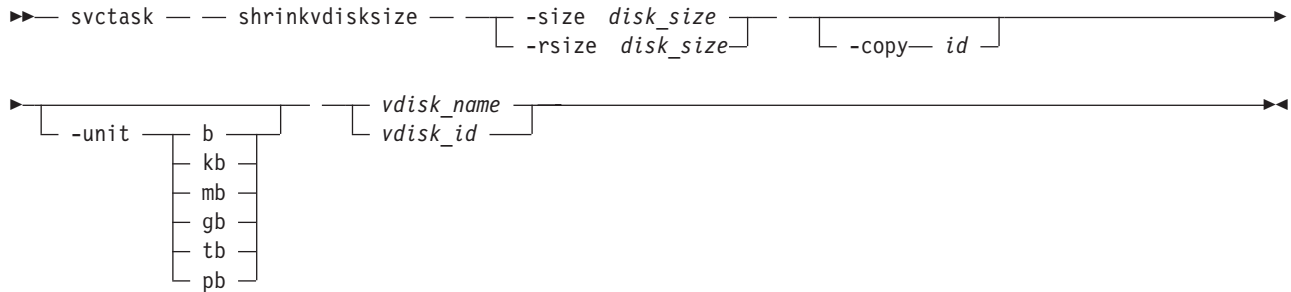
The resulting output

No feedback

shrinkvdisksize

The `shrinkvdisksize` command reduces the size of a VDisk by the specified capacity.

Syntax



Parameters

-size *disk_size*

(Required) Specifies the size reduction for the designated virtual disk. The **-size** parameter cannot be used with the **-rsize** parameter. You must specify either **-size** or **-rsize**.

-rsize *disk_size*

(Optional) Reduces the real size of a space-efficient VDisk by the specified amount. Specify the *disk_size* value using an integer. Specify the units for a *disk_size* integer using the **-unit** parameter; the default is MB. The **-rsize** value can be greater than, equal to, or less than the size of the VDisk. You must specify either the **-size** parameter or the **-rsize** parameter.

-copy *id*

(Optional) Specifies the copy to change the real capacity for. You must also specify the **-rsize** parameter. If the **-copy** parameter is not specified, all copies of the VDisk are reduced. This parameter is required if the VDisk is mirrored and only one copy is space-efficient.

-unit **b** | **kb** | **mb** | **gb** | **tb** | **pb**

(Optional) Specifies the data units to be used in conjunction with the value that is specified by the **-size** parameter.

vdisk_name | *vdisk_id*

(Required) Specifies the virtual disk that you want to modify, either by ID or by name.

Description

The `shrinkvdisksize` command reduces the capacity that is allocated to the particular virtual disk by the amount that you specify. You cannot shrink the real size of a space-efficient volume below its used size. All capacities, including changes, must be in multiples of 512 bytes. An entire extent is reserved even if it is only partially used. The default capacity units are MB.

The command can be used to shrink the physical capacity that is allocated to a particular VDisk by the specified amount. The command can also be used to shrink the virtual capacity of a space-efficient VDisk without altering the physical capacity assigned to the VDisk. To change the capacity of a non-space-efficient disk, use the **-size** parameter. To change the real capacity of a space-efficient disk, use the **-rsize** parameter. To change the virtual capacity of a space-efficient disk, use the **-size** parameter.

VDisks can be reduced in size, if required.

When the virtual size of a space-efficient VDisk is changed, the warning threshold is automatically scaled to match. The new threshold is stored as a percentage.

To run the **shrinkvdisksize** command on a mirrored VDisk, all copies of the VDisk must be synchronized.

Attention: If the VDisk contains data that is being used, do not shrink the VDisk without backing up the data first.

The cluster arbitrarily reduces the capacity of the VDisk by removing a partial, one or more extents from those allocated to the VDisk. You cannot control which extents are removed and so you cannot assume that it is unused space that is removed.

Attention:

1. If the virtual disk contains data, do not shrink the disk.
2. Some operating systems or file systems use what they consider to be the outer edge of the disk for performance reasons. This command can shrink FlashCopy target virtual disks to the same capacity as the source.
3. Before you shrink a VDisk, validate that the VDisk is not mapped to any host objects. If the VDisk is mapped, data is displayed. You can determine the exact capacity of the source or master VDisk by issuing the **svcinfolsvdisk -bytes vdiskname** command. Shrink the VDisk by the required amount by issuing the **svctask shrinkvdisksize -size disk_size -unit b | kb | mb | gb | tb | pb vdisk_name | vdisk_id** command.

An invocation example

To decrease the capacity of vdisk1 by 2 KB, enter:

```
svctask shrinkvdisksize -size 2048 -unit b vdisk1
```

The resulting output

No feedback

An invocation example

To decrease the capacity of vdisk2 by 100 MB, enter:

```
svctask shrinkvdisksize -size 100 -unit mb vdisk2
```

The resulting output

No feedback

An invocation example

To decrease the real capacity of space-efficient VDisk vdisk3 by 100 MB without changing its virtual capacity, enter:

```
svctask shrinkvdisksize -rsize 100 -unit mb vdisk3
```

The resulting output

No feedback

An invocation example

To decrease the real capacity of space-efficient VDisk copy ID 1 of mirrored VDisk vdisk4 by 100 MB, enter:

```
svctask shrinkvdisksize -rsize 100 -unit mb -copy 1 vdisk4
```

The resulting output

No feedback

An invocation example

To decrease the virtual capacity of space-efficient VDisk `vdisk5` by 1 GB without changing its real capacity, enter:

```
svctask shrinkvdisksize -size 1 -unit gb vdisk5
```

The resulting output

No feedback

splitvdiskcopy

The `splitvdiskcopy` command creates a separate VDisk from a synchronized copy of a mirrored VDisk.

Syntax

```
svctask splitvdiskcopy -copy id [-iogrp io_group_id | io_group_name]
[-node node_id | node_name] [-name new_name] [-cache readwrite | none]
[-udid udid] [-force] [vdisk_name | vdisk_id]
```

Parameters

-copy *id*

(Required) Specifies the ID of the copy to split.

-iogrp *io_group_id* | *io_group_name*

(Optional) Specifies the I/O group to add the new virtual disk to. The default is the I/O group of the specified VDisk.

-node *node_id* | *node_name*

(Optional) Specifies the preferred node ID or the name for I/O operations to this virtual disk. You can use the **-node** parameter to specify the preferred access node.

-name *new_name*

(Optional) Assigns a name to the new virtual disk.

-cache *readwrite* | *none*

(Optional) Specifies the caching options for the new virtual disk. Enter **readwrite** or **none**; the default is **readwrite**.

-udid *udid*

(Optional) Specifies the *udid* for the new VDisk. The *udid* is a required identifier for OpenVMS hosts; no other hosts use this parameter. Supported values are a decimal number 0 - 32 767, or a hexadecimal number 0 - 0x7FFF. A hexadecimal number must be preceded by **0x**; for example, **0x1234**. The default *udid* value is **0**.

-force

(Optional) Allows the split to proceed even when the specified copy is not synchronized, or even when the cache flush is likely to fail. The newly created VDisk might not be consistent.

Description

The `splitvdiskcopy` command creates a new VDisk in the specified I/O Group from a copy of the specified VDisk. If the copy that you are splitting is not synchronized, you must use the **-force** parameter. The command fails if you are attempting to remove the only synchronized copy. To avoid this, wait for

the copy to synchronize or split the unsynchronized copy from the VDisk by using the **-force** parameter. You can run the command when either VDisk copy is offline.

An invocation example

```
svctask splitvdiskcopy -copy 1 vdisk8
```

The resulting output

```
Virtual Disk, id [1], successfully created.
```

Chapter 12. Managed disk group commands

The following commands enable you to work with managed disk group options with the SAN Volume Controller.

addmdisk

The **addmdisk** command adds one or more managed disks to an existing managed disk group.

Syntax

```
| ▶▶ svctask -- addmdisk -- -mdisk [ mdisk_id_list | mdisk_name_list ] [-tier [ generic_ssd | generic_hdd ] ]
| ▶ [ mdisk_group_id | mdisk_group_name ]
```

Parameters

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Required) Specifies one or more managed disk IDs or names to add to the group.

mdisk_group_id | *mdisk_group_name*

(Required) Specifies the ID or name of the managed disk group to add the disks to. When an MDisk is added, the warning threshold for the MDisk group is automatically scaled.

-tier

Specifies the tier of the MDisk or MDisks being added.

Unless otherwise specified, the current tier value associated with the MDisk will be retained. The default value for a newly discovered unmanaged MDisk is *generic_hdd*. You can change this value by using the **chmdisk** command.

External SSDs cannot be detected automatically. If you want external SSDs to be known by the system, you must either specify the tier when adding the managed disk to the mdisk group, or use the **chmdisk** command.

Description

This command adds the managed disks that you specify to the group. The disks can be specified in terms of the managed disk ID or the managed disk name.

The managed disks must be in unmanaged mode. Disks that already belong to a group cannot be added to another group until they have been deleted from their current group. You can delete a managed disk from a group under the following circumstances:

- If the managed disk does not contain any extents in use by a virtual disk
- If you can first migrate the extents in use onto other free extents within the group.

An invocation example

```
| svctask addmdisk -mdisk mdisk13:mdisk14 -tier generic_ssd Group0
```

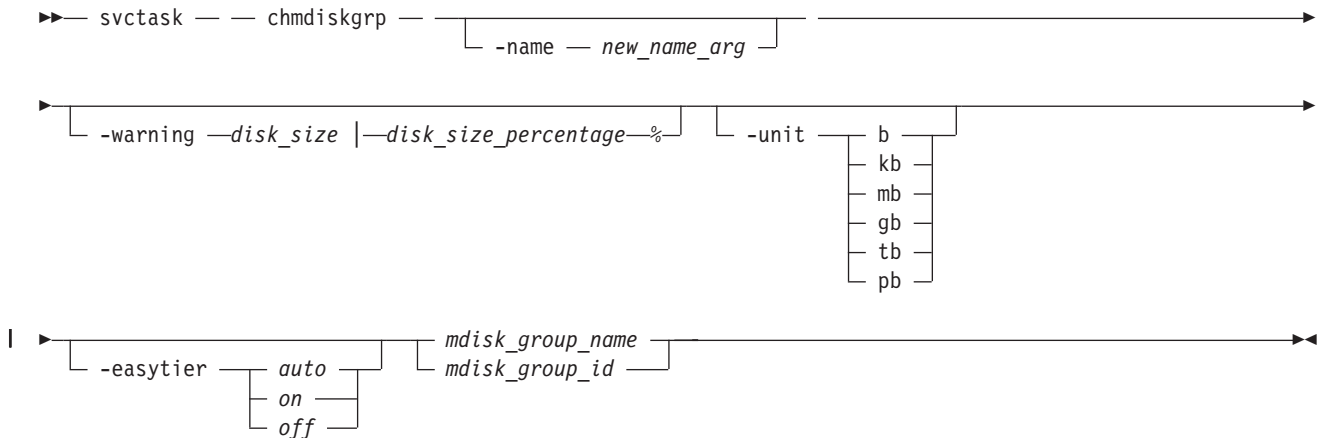
The resulting output

No feedback

chmdiskgrp

Use the **chmdiskgrp** command to modify the name that is assigned to a managed disk (MDisk) group or to set the warning threshold for the MDisk group.

Syntax



Parameters

-name *new_name_arg*

(Optional) Specifies the new name of the managed disk group.

-warning *disk_size* | *disk_size_percentage%*

(Optional) Sets a threshold at which a warning is generated. The warning is generated the first time that the threshold is exceeded by the used-disk capacity in the MDisk group. You can specify a *disk_size* integer, which defaults to megabytes (MB) unless the **-unit** parameter is specified; or you can specify a *disk_size%*, which is a percentage of the MDisk group size. To disable warnings, specify **0** or **0%**.

-unit **b** | **kb** | **mb** | **gb** | **tb** | **pb**

(Optional) Specifies the data units for the **-warning** parameter.

mdisk_group_id | *mdisk_group_name*

(Required) Specifies the ID or name of the managed disk group to modify.

-easytier

Specifies if the Easy Tier function is on or off for this MDisk group, or if it is automatically determined.

Note: **-easytier** must be followed by one of the following:

- If **-easytier** is set to *auto*, SAN Volume Controller will automatically enable Easy Tier functions when the MDisk group contains MDisks from more than one tier, and will disable automatic data placement when the MDisk group contains MDisks from only one tier.
- If **-easytier** is set to *on*, then Easy Tier functions will be active.
- If **-easytier** is set to *off*, then Easy Tier functions will be inactive.

Description

This command modifies the name, or label, assigned to a given managed disk group. Subsequently, you can use the new name to refer to the managed disk group.

The command can also be used to set the warning threshold for the managed disk group. The warning threshold is the threshold at which a warning is generated when it is exceeded by the used-disk capacity in the MDisk group.

An invocation example

```
| svctask chmdiskgrp -name testmdiskgrp -easytier on Group0
```

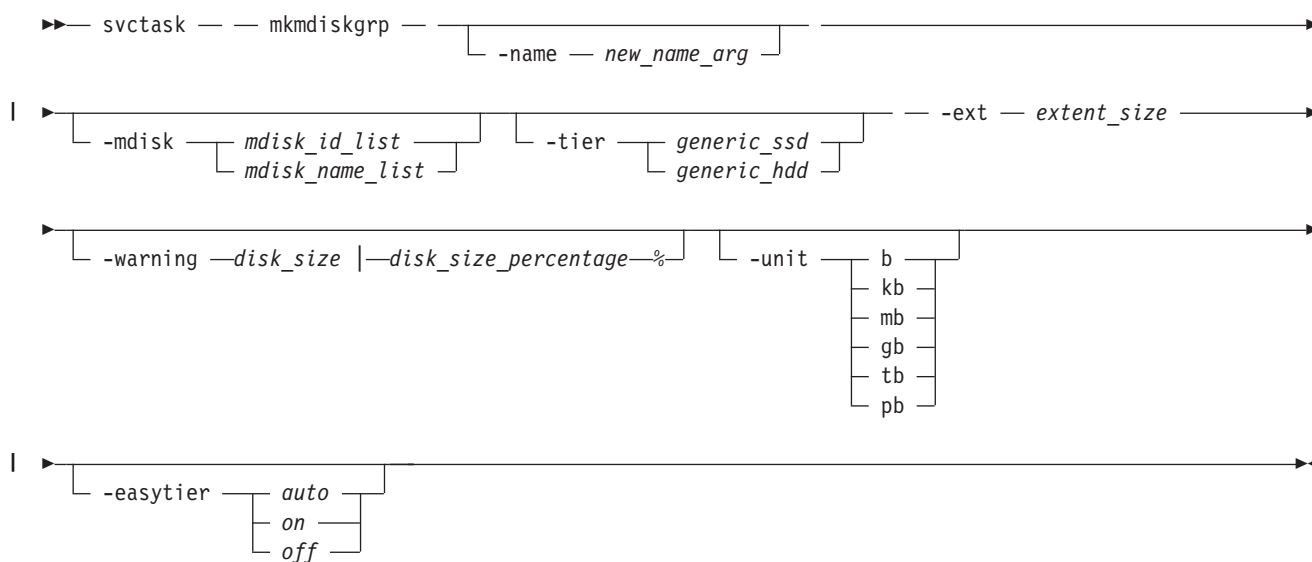
The resulting output

No feedback

mkmdiskgrp

| The **mkmdiskgrp** command creates a new managed disk group (storage pool).

Syntax



Parameters

-name *new_name_arg*

(Optional) Specifies a name to assign to the new group.

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Optional) Specifies a colon-separated list of managed disk IDs or names to add to the group. You can create an empty MDisk group by not specifying the **-mdisk** parameter.

-ext *extent_size*

| (Required) Specifies the size of the extents for this group in MB. The **ext** parameter must have one of
| the following values: **16, 32, 64, 128, 256, 512, 1024, 2048, 4096, or 8192** (MB).

-warning *disk_size* | *disk_size_percentage%*

(Optional) Generates a warning when the used disk capacity in the MDisk group first exceeds the specified threshold. You can specify a *disk_size* integer, which defaults to megabytes (MB) unless the **-unit** parameter is specified; or you can specify a *disk_size%*, which is a percentage of the MDisk group size. To disable warnings, specify **0** or **0%**. The default value is **0**.

-unit **b** | **kb** | **mb** | **gb** | **tb** | **pb**

(Optional) Specifies the data units for the **-warning** parameter.

| **-tier**

| Specifies the tier of the MDisk or MDisks being added. If you do not specify a tier, the current tier value of the MDisk will be retained. The default value for an external MDisk is `generic_hdd`.

| **-easytier**

| Specifies if the Easy Tier function is active for this MDisk group, or if it is automatically determined. `Auto` is the default value.

| **Note:**

- | • If **-easytier** is set to `auto`, SAN Volume Controller will automatically enable Easy Tier functions when the MDisk group contains MDisk from more than one tier, and will disable Easy Tier functions when the MDisk group contains MDisk from only one tier.
- | • If **-easytier** is set to `on`, then Easy Tier functions will be active.
- | • If **-easytier** is set to `off`, then Easy Tier functions will be inactive.

| **Description**

The **mkmdiskgrp** command creates a new managed disk group and assigns the group name if specified. The ID of the new group is returned if the command is successful. Managed disk groups are collections of managed disks. Each group is divided into chunks, called extents, which are used to create VDIs(volumes).

Optionally, you can specify a list of managed disks that will be added to this group. These managed disks cannot belong to another group, and they must have a mode of unmanaged. Use the **svcinfolismdiskcandidate** command to get a list of suitable candidates. If **-tier** is specified, it will apply to all of the MDisks.

Each managed disk that is a member of this group is split into extents. The storage that is available on these disks is added to a pool of extents that is available in this group. When a virtual disk is created from this group, free extents from the pool are used, in accordance with the policy used when the virtual disk was first created.

All managed disks subsequently added to this group are split into extents of the same size as the size that is assigned to the group.

When choosing an extent size, take into account the amount of storage you want to virtualize in this group. The system maintains a mapping of extents between virtual disks and managed disks. The cluster can only manage a finite number of extents (4 194 304). One cluster can virtualize the following number of extents:

- | • 64 TB – if all managed disk groups have extent sizes of 16 MB.
- | • 32 PB – if all managed disk groups have extent sizes of 8192 MB.

| **Important:** The extent size for the MDisk group can also limit VDisk size. Consider the maximum VDisk size you want to use when creating MDisk groups. Refer to the information on creating MDisk groups for a comparison of the maximum VDisk capacity for each extent size. The maximum is different for space-efficient Vdisks (thin-provisioned volumes).

Note: When an image mode VDisk is created, the MDisk group increases in capacity by the size of the image mode VDisk (not the MDisk capacity), because the image mode VDisk might be smaller than the MDisk itself. If an extent is migrated from the image mode VDisk or MDisk to elsewhere in the group, the VDisk becomes a striped VDisk (no longer image mode). At this point the available capacity might increase, because the extra capacity available on the MDisk (for example, the capacity that was not part of the image mode VDisk) becomes available.

An invocation example


```
| svctask mkmdiskgrp -mdisk mdisk13 -tier generic_hdd -easytier off -ext 512
```

The resulting output

```
MDisk Group, id [1], successfully created
```

An invocation example

```
svctask mkmdiskgrp -mdisk mdisk0:mdisk1:mdisk2:mdisk3 -ext 32
```

The resulting output

```
MDisk Group, id [0], successfully created
```

rmmdisk

The **rmmdisk** command deletes a managed disk (MDisk) from a managed disk group.

Syntax

```
▶▶ svctask — — rmmdisk — — -mdisk — [ mdisk_id_list ] — [ mdisk_name_list ] — [ -force ] —————▶▶
▶ [ mdisk_group_id ] —————▶▶
  [ mdisk_group_name ]
```

Parameters

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Required) Specifies one or more managed disk IDs or names to delete from the group.

-force

(Optional) Migrates data on the specified disks to other disks in the group. The command completes asynchronously if **-force** is specified.

mdisk_group_id | *mdisk_group_name*

(Required) Specifies the ID or name of the managed disk group to delete the disks from. The warning threshold for an MDisk group is automatically scaled when MDisks are deleted.

Description

This command attempts to remove the managed disk or disks from the group.

Deleting a managed disk from a group can only be done if the managed disk does not contain any extents in use by a virtual disk. If there are extents in use and you do not supply the force flag, the command fails.

Attention: If this disk being removed has already been powered down, removed, or is experiencing a power outage, the migration is pending and does not complete until the MDisk comes back online. The MDisk is not removed from the list of MDisks that are contained in the group.

If the disk has been deliberately removed, the only method of removing the MDisk is to remove the entire group itself.

Ensure that you do not destroy any controller LUNs until you have deleted them from the MDisk group that they belong to.

The **rmmdisk** command fails if there are insufficient free extents on other disks in the mdisk group for the duration of the command.

If you do specify the force flag, an attempt will be made to migrate the extents that are in use onto other free extents within the group. If there are not enough free extents in the group, the command will fail even if the force flag is specified.

- | When an array MDisk is in a storage pool, five extents in the storage pool are reserved for internal use. If you attempt to remove an MDisk when an array MDisk is in the storage pool, the command will fail (even if the **-force** flag is specified), if five free extents do not remain in the storage pool.

To delete the disks from the group, you have the following options:

- You can delete the virtual disk that is using the extents specified on the managed disk.
- You can add more managed disks to the group, rerun the command and specify the **-force** parameter.

When data is being migrated from the managed disk, it might take some time for the command to complete. The command itself will return with a success code, notifying you that migration is in progress. An event is logged when the migration is complete and the disk is deleted from the group at this time. You can also check the progress of any active migrations by running the **svcinfolsmigrate** command.

- | If the **-force** parameter is used, the **rmmdisk** command fails if offline Managed Disks or no online quorum disks will prevent the migration. Correct the offline or quorum disk condition and try reissuing the command.

An invocation example

```
svctask rmmdisk -mdisk mdisk12 -force Group3
```

The resulting output

No feedback

rmmdiskgrp

The **rmmdiskgrp** command deletes a managed disk group so that there is no possibility to recover it.

Syntax

```
▶▶ svctask — — rmmdiskgrp — [ -force ] [ mdisk_group_id | mdisk_group_name ] ▶▶
```

Parameters

-force

(Optional) Specifies that all virtual disks and virtual disk-to-host mappings be deleted.

Attention: Use this parameter with extreme caution. When you use this parameter, all managed disks in the group are removed and the group itself is deleted.

mdisk_group_id | *mdisk_group_name*

(Required) Specifies the ID or name of the managed disk group that is to be deleted.

Description

The **rmmdiskgrp** command deletes the specified managed disk group. The **-force** parameter is required if there are virtual disks that have been created from this group or if there are managed disks in the group. Otherwise, the command fails.

Deleting a managed disk group is essentially the same as deleting a cluster or part of a cluster, because the managed disk group is the central point of control of virtualization. Because virtual disks are created using available extents in the group, mapping between virtual disk extents and managed disk extents is controlled based on the group.

The command deletes all VDisk copies in the specified MDisk group. If the VDisk has no remaining synchronized copies in other MDisk groups, the VDisk is also deleted.

Attention:

1. This command partially completes asynchronously. All virtual disks, host mappings, and Copy Services relationships are deleted before the command completes. The deletion of the managed disk group then completes asynchronously.
2. Before you issue the command, ensure that you want to delete all mapping information; data that is contained on virtual disks cannot be recovered after the managed disk group has been deleted.

In detail, if you specify the **-force** parameter and the virtual disks are still using extents in this group, the following actions are initiated or occur:

- The mappings between that disk and any host objects and the associated Copy Services relationships are deleted.
- If the virtual disk is a part of a FlashCopy mapping, the mapping is deleted.

Note: If the mapping is not in the **idle_or_copied** or **stopped** states, the mapping is force-stopped and then deleted. Force-stopping the mapping might cause other FlashCopy mappings in the cluster to also be stopped. See the description for the **-force** parameter in the **stopfcmap** command for additional information.

- Any virtual disk that is in the process of being migrated into or out of the managed disk group is deleted. This frees up any extents that the virtual disk was using in another managed disk group.
- Virtual disks are deleted without first flushing the cache. Therefore, the storage controller LUNs that underlie any image mode MDisks might not contain the same data as the image mode VDisk prior to the deletion.
- If there are managed disks in the group, all disks are deleted from the group. They are returned to the unmanaged state.
- The group is deleted.

Attention: If you use the **-force** parameter to delete all the managed disk groups in your cluster, you are returned to the processing state where you were after you added nodes to the cluster. All data that is contained on the virtual disks is lost and cannot be recovered.

An invocation example

```
svctask rmdiskgrp -force Group3
```

The resulting output

```
No feedback
```

Chapter 13. Managed disk commands

The following commands enable you to work with managed disk options with the SAN Volume Controller.

If the cluster detects an MDisk, it automatically adds it to the list of known MDisks. If you subsequently delete the RAID that corresponds to the MDisk, the cluster only deletes the MDisk from the list if the MDisk is offline and it has a mode of unmanaged (it does not belong to an MDisk group).

applymdisksoftware (Discontinued)

| **Attention:** The **applymdisksoftware** command has been discontinued. Use the **applydrivesoftware** command to upgrade drives.

chmdisk

Use the **chmdisk** command to modify the name of a managed disk (MDisk).

Syntax

```
| ▶▶ svctask -- chmdisk -- [-name new_name_arg] [-tier generic_ssd | generic_hdd]
```



```
| ▶ [mdisk_id | mdisk_name]
```

Parameters

-name *new_name_arg*
Specifies the new name to be applied to the managed disk.

| **-tier**
Specifies the new tier of the MDisk.

mdisk_id | *mdisk_name*
(Required) Specifies the ID or name of the managed disk to modify.

Description

| This command modifies the attributes of a managed disk.

An invocation example

| `svctask chmdisk -tier generic_hdd mdisk13`

The resulting output

No feedback

chquorum

| Use the **chquorum** command to change the quorum association.

Syntax

```

| ►► svctask — — chquorum — — [ -active ] — — [ -mdisk —mdisk_id_or_name— | -drive —drive_id— ] — — quorum_id —►►

```

Parameters

-active

(Optional) Makes the specified quorum ID the active one. The **active** parameter must be used if neither the **mdisk** nor the **drive** parameters are specified.

-mdisk *mdisk_id_or_name* | -drive *drive_id*

(Optional) Specifies the MDisk or drive to be this quorum ID.

Note: SAN Volume Controller uses MDisks only.

quorum_id

(Required) Specifies which quorum ID to change.

Description

Use the **chquorum** command to change the quorum association. To identify the drive or MDisk that is the current active quorum disk, use the **svcinfolquorum** command.

Note: SAN Volume Controller uses MDisks only.

Attention: Only assign quorum to drives in the control enclosure or to external MDisks. Some maintenance procedures require that quorum is moved temporarily to expansion enclosures. Once that procedure is complete, return the quorum drives to the control enclosure.

The **chquorum** command is not synchronous, but usually takes only a few seconds to complete. In some situations it can take several minutes.

The cluster uses the quorum disk or drive as a tie breaker when exactly half of the nodes that were previously a member of the cluster is present.

The use of a quorum disk or drive allows the cluster to manage a SAN fault that splits the cluster exactly in half. One half of the cluster continues to operate and the other half stops until SAN connectivity is restored.

There is only one quorum disk or drive; however, the cluster uses three as quorum candidates. The cluster selects the actual quorum disk or drive from the pool of quorum candidates. The quorum candidates also hold a copy of important cluster metadata. Just over 256 MB is reserved for this purpose on each quorum candidate disk. When using an MDisk as quorum disk, this space is allocated from the MDisk group (storage pool). The number of extents required depends on the extent size for the managed disk group containing the MDisk. Table 15 provides the number of extents reserved for quorum use by extent size.

Table 15. Number of extents reserved by extent size

Extent size (MB)	Number of extents reserved for quorum use
16	17
32	9
64	5
128	3

Table 15. Number of extents reserved by extent size (continued)

Extent size (MB)	Number of extents reserved for quorum use
256	2
512	1
1024	1
2048	1
4096	1
8192	1

When you issue this command, the MDisk or drive that currently is assigned the quorum index number is set to a nonquorum disk. The cluster automatically assigns quorum indexes.

You can set the active quorum disk or drive with the **active** parameter. This can be useful in a split-site cluster configuration to ensure that the most highly-available quorum disk or drive is used.

An invocation example

```
svctask chquorum -mdisk 45 3
```

The resulting output

No feedback

dumpallmdiskbadblocks

The **dumpallmdiskbadblocks** command dumps bad block counts to a dump file for use by Directed Maintenance Procedures (DMPs) and the satask **snap** command.

Syntax

```
▶▶— svctask — — dumpallmdiskbadblocks —————▶▶
```

Parameters

None

Description

Use the svctask **dumpallmdiskbadblocks** command to dump bad block counts to a readable ASCII dump file for use by DMPs and the satask **snap** command. The output contains bad blocks for which an error log has been raised.

Use svcinfo **lsdumps -prefix /dumps/mdisk** to list the output files. Use svctask **cleardumps -prefix /dumps/mdisk** to clear the output files.

The maximum number of dump files is 20.

An invocation example

```
svctask dumpallmdiskbadblocks
```

The resulting output if MDisk 2 and MDisk 5 have bad blocks

Cluster name: my_cluster
Timestamp of dump: Fri Oct 31 11:27:33 2009 UTC

Mdisk id: 2
Mdisk name: mdisk2
Number of bad blocks: 4

Mdisk id: 5
Mdisk name: mdisk 5
Number of bad blocks: 1

Total mdisks with bad blocks: 2
Total number of bad blocks: 5

The resulting output if the MDisks have no bad blocks

Cluster name: my_cluster
Timestamp of dump: Fri Oct 31 11:27:33 2009 UTC

Total mdisks with bad blocks: 0
Total number of bad blocks: 0

dumpmdiskbadblocks

The **dumpmdiskbadblocks** command writes the bad block counts and locations that are on a specified MDisk to a dump file for use by Directed Maintenance Procedures (DMPs).

Syntax

▶— svctask — — dumpmdiskbadblocks — — *object_id_or_name* —————▶

Parameters

object_id_or_name

(Required) Specifies the MDisk for which you need to dump the bad block record table.

Description

Use the svctask **dumpmdiskbadblocks** command to write the bad block counts and locations that are on a specified MDisk to a readable ASCII dump file for use by DMPs. The output consists of bad blocks for which an error log has been raised.

Use svcinfo **lsdumps -prefix /dumps/mdisk** to list the output files. Use svctask **cleardumps -prefix /dumps/mdisk** to clear the output files.

The reported error log sequence numbers correspond to the first error seen in the bad block record, which is a 512-block region.

- If there are multiple error logs in the same region, the earliest error sequence is used.
- If there are error logs of different types in the same region, error sequence numbers for killsectors take precedence.
- If a range of bad blocks runs across record boundaries, the sequence number corresponding to the last record is used.

The maximum number of dump files is 20.

An invocation example

```
svctask dumpmdiskbadblocks 3
```


The resulting output if the MDisk has bad blocks

```
Cluster name: my_cluster
Timestamp of dump: Fri Oct 31 11:27:33 2009 UTC
```

```
Mdisk id: 3
Mdisk name: mdisk3
Number of bad blocks: 6
```

```
Start LBA: 0x1234123412341234
Length: 2
Error log sequence number: 1
```

```
Start LBA: 0x5678567812341234
Length: 4
Error log sequence number: 2
```

The resulting output if the MDisk has no bad blocks

```
Cluster name: my_cluster
Timestamp of dump: Fri Oct 31 11:27:33 2009 UTC
```

```
Mdisk id: 3
Mdisk name: mdisk3
Number of bad blocks: 0
```

includemdisk

Use the **includemdisk** command to include a disk that has been excluded by the cluster.

Syntax

```
svctask — — includemdisk — — mdisk_id —————▶
                               └── mdisk_name ───┘
```

Parameters

mdisk_id | *mdisk_name*

(Required) Specifies the ID or name of the managed disk to add back into the cluster.

Description

The specified managed disk is included in the cluster.

You might exclude a disk from the cluster because of multiple I/O failures. These failures might be caused by noisy links. Once a fabric-related problem has been fixed, the excluded disk can be added back into the cluster.

Running this command against an MDisk might change its state, whether the state is reported as excluded.

Note: If an MDisk is in the excluded state, is offline, and does not belong to an MDisk group, issuing an include command for this MDisk results in the MDisk record being deleted from the cluster.

An invocation example

```
svctask includemdisk mdisk5
```

The resulting output

```
No feedback
```

setquorum (Deprecated)

- | **Attention:** The **svctask setquorum** command is deprecated. Use the **svctask chquorum** command to change the quorum association.

triggermdiskdump (Discontinued)

- | **Attention:** The **triggermdiskdump** command is discontinued. Use the **triggerdrivedump** command to collect support data from a disk drive.

Chapter 14. IBM FlashCopy commands

The following commands enable you to work with FlashCopy methods and functions with the SAN Volume Controller.

chfconsistgrp

The **chfconsistgrp** command changes the name of a consistency group or marks the group for auto-deletion.

Syntax

```
svctask -- chfconsistgrp -- [-name new_name_arg]
                             [ -autodelete on | off ] [ fc_consist_group_id | fc_consist_group_name ]
```

Parameters

-name *new_name_arg*

(Optional) Specifies the new name to assign to the consistency group.

-autodelete on | off

(Optional) Deletes the consistency group when the last mapping that it contains is deleted or removed from the consistency group.

fc_consist_group_id | *fc_consist_group_name*

(Required) Specifies the ID or existing name of the consistency group that you want to modify.

Description

The **chfconsistgrp** command changes the name of a consistency group, marks the group for auto-deletion, or both.

An invocation example

```
svctask chfconsistgrp -name testgrp1 fcconsistgrp1
```

The resulting output


No feedback

chfcmap

The **chfcmap** command modifies attributes of an existing mapping.

Syntax

```
svctask -- chfcmap -- [-name new_name_arg] [-force]
```



Parameters

-name *new_name_arg*

(Optional) Specifies the new name to assign to the mapping. The **-name** parameter cannot be used with any other optional parameters.

-force

(Optional) Specifies that the mapping be modified to a stand-alone mapping (equivalent to creating the mapping without a consistency group ID). You cannot specify the **-force** parameter with the **-consistgrp** parameter.

-consistgrp *consist_group_id* | *consist_group_name*

(Optional) Specifies the consistency group for which you want to modify the mapping. You cannot specify the **-consistgrp** parameter with the **-force** parameter.

Note: The consistency group cannot be modified if the specified consistency group is in the **preparing**, **prepared**, **copying**, **suspended**, or **stopping** state.

-copyrate *rate*

(Optional) Specifies the copy rate. The *rate* value can be 0 - 100. The default value is 50. A value of 0 indicates no background copy process. For the supported **-copyrate** values and their corresponding rates, see Table 16 on page 195.

-autodelete **on** | **off**

(Optional) Specifies that the autodelete function be turned on or off for the specified mapping. When you specify the **-autodelete on** parameter, you are deleting a mapping after the background copy completes. If the background copy is already complete, the mapping is deleted immediately.

-cleanrate *rate*

(Optional) Sets the cleaning rate for the mapping. The *rate* value can be 0 - 100. The default value is 50.

fc_map_id | *fc_map_name*

(Required) Specifies the ID or name of the mapping to modify. Enter the ID or name last on the command line.

Description

The **svctask chfcmap** command modifies attributes of an existing mapping.

Attention: You must enter the *fc_map_id* | *fc_map_name* last on the command line.

If you have created several FlashCopy mappings for a group of VDisks that contain elements of data for the same application, you can assign these mappings to a single FlashCopy consistency group. You can then issue a single prepare command and a single start command for the whole group, for example, so that all of the files for a particular database are copied at the same time.

The **copyrate** parameter specifies the copy rate. If 0 is specified, background copy is disabled. The **cleanrate** parameter specifies the rate for cleaning the target VDisk. The cleaning process is only active if the mapping is in the **copying** state and the background copy has completed, the mapping is in the **copying** state and the background copy is disabled, or the mapping is in the **stopping** state. You can

disable cleaning when the mapping is in the **copying** state by setting the **cleanrate** parameter to 0. If the **cleanrate** is set to 0, the cleaning process runs at the default rate of 50 when the mapping is in the **stopping** state to ensure that the stop operation completes.

Table 16 provides the relationship of the copy *rate* and cleaning *rate* values to the attempted number of grains to be split per second. A grain is the unit of data represented by a single bit.

Table 16. Relationship between the rate, data rate and grains per second values

User-specified <i>rate</i> attribute value	Data copied/sec	256 KB grains/sec	64 KB grains/sec
1 - 10	128 KB	0.5	2
11 - 20	256 KB	1	4
21 - 30	512 KB	2	8
31 - 40	1 MB	4	16
41 - 50	2 MB	8	32
51 - 60	4 MB	16	64
61 - 70	8 MB	32	128
71 - 80	16 MB	64	256
81 - 90	32 MB	128	512
91 - 100	64 MB	256	1024

An invocation example

```
svctask chfcmap -name testmap 1
```

The resulting output

No feedback

mkfcconsistgrp

The **mkfcconsistgrp** command creates a new FlashCopy consistency group and identification name.

Syntax

```

▶▶ svctask — — mkfcconsistgrp — [ -name — consist_group_name ] [ -autodelete ] ▶▶

```

Parameters

-name *consist_group_name*

(Optional) Specifies a name for the consistency group. If you do not specify a consistency group name, a name is automatically assigned to the consistency group. For example, if the next available consistency group ID is id=2, the consistency group name is fccstgrp2.

-autodelete

(Optional) Deletes the consistency group when the last mapping that it contains is deleted or removed from the consistency group.

Description

This command creates a new consistency group and identification name. The ID of the new group is displayed when the command process completes.

If you have created several FlashCopy mappings for a group of VDisks that contain elements of data for the same application, you might find it convenient to assign these mappings to a single FlashCopy consistency group. You can then issue a single prepare command and a single start command for the whole group, for example, so that all of the files for a particular database are copied at the same time.

An invocation example

```
svctask mkfcconsistgrp
```

The resulting output

```
FlashCopy Consistency Group, id [1], successfully created
```

mkfcmap

The **mkfcmap** command creates a new FlashCopy mapping, which maps a source VDisk (volume) to a target VDisk (volume) for subsequent copying.

Syntax

```
svctask -- mkfcmap -- -source [ src_vdisk_id | src_vdisk_name ]
-- -target [ target_vdisk_id | target_vdisk_name ] [-name new_name_arg]
-- [-consistgrp [ consist_group_id | consist_group_name ] [-copyrate rate] [-autodelete]
-- [-grainsize [ 64 | 256 ] [-incremental] [-cleanrate rate]
-- [-iogrp [ iogroup_name | iogroup_id ]
```

Parameters

-source *src_vdisk_id* | *src_vdisk_name*

(Required) Specifies the ID or name of the source VDisk (volume).

-target *target_vdisk_id* | *target_vdisk_name*

(Required) Specifies the ID or name of the target VDisk (volume).

-name *new_name_arg*

(Optional) Specifies the name to assign to the new mapping.

-consistgrp *consist_group_id* | *consist_group_name*

(Optional) Specifies the consistency group to add the new mapping to. If you do not specify a consistency group, the mapping is treated as a stand-alone mapping.

-copyrate *rate*

(Optional) Specifies the copy rate. The *rate* value can be 0 - 100. The default value is 50. A value of 0 indicates no background copy process. For the supported **-copyrate** values and their corresponding rates, see Table 17 on page 198.

-autodelete

(Optional) Specifies that a mapping be deleted when the background copy completes. The default, which applies if this parameter is not entered, is that **autodelete** is set to off.

-grainsize 64 | 256

(Optional) Specifies the grain size for the mapping. The default value is 256. Once set, this value cannot be changed.

-incremental

(Optional) Marks the FlashCopy mapping as an incremental copy. The default is nonincremental. Once set, this value cannot be changed.

-cleanrate *rate*

(Optional) Sets the cleaning rate for the mapping. The *rate* value can be 0 - 100. The default value is 50.

-iogrp *iogroup_name* | *iogroup_id*

(Optional) Specifies the I/O group for the FlashCopy bitmap. Once set, this value cannot be changed. The default I/O group is either the source VDisk (volume), if a single target map, or the I/O group of the other FlashCopy mapping to which either the source or target VDIsks (volumes) belong.

Note: If not enough bitmap space is available to complete this command, more space will automatically be allocated in the bitmap memory (unless you have already reached the maximum bitmap memory).

Description

This command creates a new FlashCopy mapping. This mapping persists until it is manually deleted, or until it is automatically deleted when the background copy completes and the **autodelete** parameter set to **on**. The source and target VDIsks (volumes) must be specified on the **mkfcmap** command. The **mkfcmap** command fails if the source and target VDIsks (volumes) are not identical in size. Issue the **svcinfolsvdisk -bytes** command to find the exact size of the source VDisk (volume) for which you want to create a target disk of the same size. The target VDisk (volume) that you specify cannot be a target VDisk (volume) in an existing FlashCopy mapping. A mapping cannot be created if the resulting set of connected mappings exceeds 256 connected mappings.

The mapping can optionally be given a name and assigned to a consistency group, which is a group of mappings that can be started with a single command. These are groups of mappings that can be processed at the same time. This enables multiple VDIsks (volumes) to be copied at the same time, which creates a consistent copy of multiple disks. This consistent copy of multiple disks is required by some database products in which the database and log files reside on different disks.

If the specified source and target VDIsks (volumes) are the target and source VDIsks (volumes), respectively, of an existing mapping, then the mapping being created and the existing mapping become partners. If one mapping is created as incremental, then its partner is automatically incremental. A mapping can have only one partner.

The **copyrate** parameter specifies the copy rate. If 0 is specified, background copy is disabled. The **cleanrate** parameter specifies the rate for cleaning the target VDisk (volume). The cleaning process is only active if the mapping is in the copying state and the background copy has completed, the mapping is in the copying state and the background copy is disabled, or the mapping is in the stopping state. You can disable cleaning when the mapping is in the copying state by setting the **cleanrate** parameter to 0. If the **cleanrate** is set to 0, the cleaning process runs at the default rate of 50 when the mapping is in the stopping state to ensure that the stop operation completes.

Table 17 on page 198 provides the relationship of the copy rate and cleaning rate values to the attempted number of grains to be split per second. A grain is the unit of data represented by a single bit.

Table 17. Relationship between the rate, data rate and grains per second values

User-specified <i>rate</i> attribute value	Data copied/sec	256 KB grains/sec	64 KB grains/sec
1 - 10	128 KB	0.5	2
11 - 20	256 KB	1	4
21 - 30	512 KB	2	8
31 - 40	1 MB	4	16
41 - 50	2 MB	8	32
51 - 60	4 MB	16	64
61 - 70	8 MB	32	128
71 - 80	16 MB	64	256
81 - 90	32 MB	128	512
91 - 100	64 MB	256	1024

An invocation example

```
svctask mkfcmap -source 0 -target 2 -name mapone
```

The resulting output

```
FlashCopy Mapping, id [1], successfully created
```

prestartfcconsistgrp

The **prestartfcconsistgrp** command prepares a consistency group (a group of FlashCopy mappings) so that the consistency group can be started. This command flushes the cache of any data that is destined for the source VDisk and forces the cache into the write-through mode until the consistency group is started.

Syntax

```

▶▶ svctask — — prestartfcconsistgrp — [ -restore ] [ fc_consist_group_id | fc_consist_group_name ]

```

Parameters

-restore

(Optional) Specifies the restore flag. This forces the consistency group to be prepared even if the target VDisk of one of the mappings in the consistency group is being used as a source VDisk of another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_consist_group_id | *fc_consist_group_name*

(Required) Specifies the name or ID of the consistency group that you want to prepare.

Description

This command prepares a consistency group (a group of FlashCopy mappings) to subsequently start. The preparation step ensures that any data that resides in the cache for the source VDisk is first flushed to disk. This step ensures that the FlashCopy target VDisk is identical to what has been acknowledged to the host operating system as having been written successfully to the source VDisk.

You can use the **restore** parameter to force the consistency group to be prepared even if the target VDisk of one or more mappings in the consistency group is being used as a source VDisk of another active

mapping. In this case the mapping restores as shown in the **svcinfolsfcmmap** view. If the **restore** parameter is specified when preparing a consistency group where none of the target VDisks are the source VDisk of another active mapping, then the parameter is ignored.

You must issue the **svctask prestartfcconsistgrp** command to prepare the FlashCopy consistency group before the copy process can be started. When you have assigned several mappings to a FlashCopy consistency group, you must issue a single prepare command for the whole group to prepare all of the mappings at once.

The consistency group must be in the `idle_or_copied` or `stopped` state before it can be prepared. When you enter the **prestartfcconsistgrp** command, the group enters the preparing state. After the preparation is complete, the consistency group status changes to `prepared`. At this point, you can start the group.

If FlashCopy mappings are assigned to a consistency group, the preparing and the subsequent starting of the mappings in the group must be performed on the consistency group rather than on an individual FlashCopy mapping that is assigned to the group. Only stand-alone mappings, which are mappings that are not assigned to a consistency group, can be prepared and started on their own. A FlashCopy mapping must be prepared before it can be started.

An invocation example

```
svctask prestartfcconsistgrp 1
```

The resulting output

No feedback

prestartfcmap

The **prestartfcmap** command prepares a FlashCopy mapping so that it can be started. This command flushes the cache of any data that is destined for the source VDisk and forces the cache into the write-through mode until the mapping is started.

Syntax

```
svctask -- prestartfcmap -- [-restore] [fc_map_id | fc_map_name]
```

Parameters

-restore

(Optional) Specifies the restore flag. This forces the mapping to be prepared even if the target VDisk is being used as a source VDisk in another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_map_id | *fc_map_name*

(Required) Specifies the name or ID of the mapping to prepare.

Description

This command prepares a single mapping for subsequent starting. The preparation step ensures that any data that resides in the cache for the source VDisk is first transferred to disk. This step ensures that the copy that is made is consistent with what the operating system expects on the disk.

The **restore** parameter can be used to force the mapping to be prepared even if the target VDisk is being used as a source VDisk of another active mapping. In this case, the mapping is restoring as shown in the **svcinfolsfcmmap** view. If the **restore** parameter is specified when preparing a mapping where the target VDisk is not the source VDisk of another active mapping, then the parameter is ignored.

Note: To prepare a FlashCopy mapping that is part of a consistency group, you must use the **prestartfcconsistgrp** command.

The mapping must be in the `idle_or_copied` or `stopped` state before it can be prepared. When the **prestartfcmap** command is processed, the mapping enters the preparing state. After the preparation is complete, it changes to the prepared state. At this point, the mapping is ready to start.

Attention: This command can take a considerable amount of time to complete.

An invocation example

```
svctask prestartfcmap 1
```

The resulting output

No feedback

rmfcconsistgrp

The **rmfcconsistgrp** command deletes a FlashCopy consistency group.

Syntax

```
▶▶ svctask — — rmfcconsistgrp — [ -force ] [ fc_consist_group_id | fc_consist_group_name ] ▶▶
```

Parameters

-force

(Optional) Specifies that all of the mappings that are associated with a consistency group that you want to delete are removed from the group and changed to stand-alone mappings. This parameter is only required if the consistency group that you want to delete contains mappings.

fc_consist_group_id | *fc_consist_group_name*

(Required) Specifies the ID or name of the consistency group that you want to delete.

Description

This command deletes the specified FlashCopy consistency group. If there are mappings that are members of the consistency group, the command fails unless you specify the **-force** parameter. When you specify the **-force** parameter, all of the mappings that are associated with the consistency group are removed from the group and changed to stand-alone mappings.

To delete a single mapping in the consistency group, you must use the **svctask rmfcmap** command.

An invocation example

```
svctask rmfcconsistgrp fcconsistgrp1
```

The resulting output

No feedback

rmfcmap

The **rmfcmap** command deletes an existing mapping.

Syntax

```
svctask -- rmfcmap -- [-force] [fc_map_id | fc_map_name]
```

Parameters

-force

(Optional) Specifies that the target VDisk is brought online. This parameter is required if the FlashCopy mapping is in the stopped state.

fc_map_id | *fc_map_name*

(Required) Specifies the ID or name of the FlashCopy mapping to delete. Enter the ID or name last on the command line.

Description

The **rmfcmap** command deletes the specified mapping if the mapping is in the `idle_or_copied` or `stopped` state. If it is in the `stopped` state, the **-force** parameter is required. If the mapping is in any other state, you must stop the mapping before you can delete it.

Deleting a mapping only deletes the logical relationship between the two virtual disks; it does not affect the virtual disks themselves. However, if you force the deletion, the target virtual disk (which might contain inconsistent data) is brought back online.

An invocation example

```
svctask rmfcmap testmap
```

The resulting output

No feedback

startfcconsistgrp

The **startfcconsistgrp** command starts a FlashCopy consistency group of mappings. This command makes a point-in-time copy of the source VDisks at the moment that the command is started.

Syntax

```
svctask -- startfcconsistgrp -- [-prep] [-restore] [fc_consist_group_id | fc_consist_group_name]
```

Parameters

-prep

(Optional) Specifies that the designated FlashCopy consistency group be prepared prior to starting the FlashCopy consistency group. A FlashCopy consistency group must be prepared before it can be started. When you use this parameter, the system automatically issues the **prestartfcconsistgrp** command for the group that you specify.

-restore

(Optional) Specifies the restore flag. When combined with the **prep** option, this forces the consistency

group to be prepared even if the target VDisk of one of the mappings in the consistency group is being used as a source VDisk in another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_consist_group_id | *fc_consist_group_name*

(Required) Specifies the ID or name of the consistency group mapping to start.

Description

This command starts a consistency group, which results in a point-in-time copy of the source VDIs of all mappings in the consistency group. You can combine the **restore** parameter with the **prep** parameter to force the consistency group to be prepared prior to starting, even if the target VDisk of one or more mappings in the consistency group is being used as a source VDisk of another active mapping. In this case, the mapping is restoring as shown in the **svcinfc lsfcmap** view. If the **restore** parameter is specified when starting a consistency group where none of the target VDIs are the source VDisk of another active mapping, then the parameter is ignored.

If a consistency group is started and the target VDisk of the mapping being started has up to four other incremental FlashCopy mappings using the target, the incremental recording is left on. If there are more than four other incremental FlashCopy mappings using the target VDisk, the incremental recording for all of these mappings is turned off until they are restarted.

Note: The **startfcconsistgrp** command can take some time to process particularly if you have specified the **prep** parameter. If you use the **prep** parameter, you give additional processing control to the system because the system must prepare the mapping before the mapping is started. If the prepare process takes too long, the system completes the prepare but does not start the consistency group. In this case, error message CMMVC6209E displays. To control the processing times of the **prestartfcconsistgrp** and **startfcconsistgrp** commands independently of each other, do not use the **prep** parameter. Instead, first issue the **prestartfcconsistgrp** command, and then to start the copy, issue the **startfcconsistgrp** command.

An invocation example

```
svctask startfcconsistgrp -prep 2
```

The resulting output

No feedback

startfcmap

The **startfcmap** command starts a FlashCopy mapping. This command makes a point-in-time copy of the source VDisk at the moment that the command is started.

Syntax

```
➤— svctask — — startfcmap — [ -prep ] [ -restore ] [ fc_map_id | fc_map_name ] —➤
```

Parameters

-prep

(Optional) Specifies that the designated mapping be prepared prior to starting the mapping. A mapping must be prepared before it can be started. When you use this parameter, the system automatically issues the **prestartfcmap** command for the group that you specify.

-restore

(Optional) Specifies the restore flag. When combined with the **prep** option, this forces the mapping to

be prepared even if the target VDisk is being used as a source VDisk in another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_map_id | fc_map_name
(Required) Specifies the ID or name of the mapping to start.

Description

This command starts a single mapping, which results in a point-in-time copy of the source VDisk. You can combine the **restore** parameter with the **prep** parameter to force the mapping to be prepared prior to starting, even if the target VDisk is being used as a source VDisk of another active mapping. In this case, the mapping is restoring as shown in the **svcinfolsfcmmap** view. If the **restore** parameter is specified when starting a mapping where the target VDisk is not the source VDisk of another active mapping, then the parameter is ignored and the mapping is not restoring as shown in the **svcinfolsfcmmap** view.

If a mapping is started and the target VDisk of the mapping being started has up to four other incremental FlashCopy mappings using the target, the incremental recording is left on. If there are more than four other incremental FlashCopy mappings using the target VDisk, the incremental recording for all of these mappings is turned off until they are restarted.

Note: The **startfcmap** command can take some time to start, particularly if you use the **prep** parameter. If you use the **prep** parameter, you give additional starting control to the system. The system must prepare the mapping before the mapping is started. To keep control when the mapping starts, you must issue the **prestartfcmap** command before you issue the **startfcmap** command.

An invocation example

```
svctask startfcmap -prep 2
```

The resulting output

No feedback

stopfcconsistgrp

The **stopfcconsistgrp** command stops all processing that is associated with a FlashCopy consistency group that is in one of the following processing states: prepared, copying, stopping, or suspended.

Syntax

```
svctask — — stopfcconsistgrp — [ -force ] [ -split ] [ fc_consist_group_id | fc_consist_group_name ]
```

Parameters

-force

(Optional) Specifies that all processing that is associated with the mappings of the designated consistency group be stopped immediately.

Note: When you use this parameter, all FlashCopy mappings that depend on the mappings in this group (as listed by the **lsfcmapdependentmaps** command) are also stopped.

-split

(Optional) Breaks the dependency on the source VDisks of any mappings that are also dependent on the target VDisk. This parameter can only be specified when stopping a consistency group where all maps in the group have progress of 100 as shown by the **svcinfolsfcmmap** command.

fc_consist_group_id | *fc_consist_group_name*

(Required) Specifies the name or ID of the consistency group that you want to stop.

Description

This command stops a group of mappings in a consistency group. If the copy process is stopped, the target disks become unusable unless they already contain complete images of the source. Disks that contain complete images of the source have a progress of 100, as indicated in the **lsfcmap** command output. The target VDisk is reported as offline if it does not contain a complete image. Before you can access this VDisk, the group of mappings must be prepared and restarted.

If the consistency group is in the *idle_or_copied* state, then the `stopfcconsistgrp` command has no effect and the consistency group stays in the *idle_or_copied* state.

Note: Prior to SVC 4.2.0, the **stopfcconsistgrp** command always caused the consistency group to go to the stopped state, taking the target VDIsks offline.

The **split** option can be used when all of the maps in the group have progress of 100. It removes the dependency of any other maps on the source VDIsks. It might be used prior to starting another FlashCopy consistency group whose target disks are the source disks of the mappings being stopped. Once the consistency group has been stopped with the **split** option, the other consistency group could then be started without the **restore** option.

An invocation example

```
svctask stopfcconsistgrp testmapone
```

The resulting output

No feedback

stopfcmap

The **stopfcmap** command stops all processing that is associated with a FlashCopy mapping that is in one of the following processing states: prepared, copying, stopping, or suspended.

Syntax

```
▶▶ svctask — — stopfcmap — [ -force ] [ -split ] [ fc_map_id ] [ fc_map_name ] ▶▶
```

Parameters

-force

(Optional) Specifies that all processing that is associated with the designated mapping be stopped immediately.

Note: When you use this parameter, all FlashCopy mappings that depend on this mapping (as listed by the **lsfcmapdependentmaps** command) are also stopped.

-split

(Optional) Breaks the dependency on the source VDisk of any mappings that are also dependent on the target disk. This parameter can only be specified when stopping a map that has progress of 100 as shown by the **svcinfolsfcmmap** command.

fc_map_id | *fc_map_name*

(Required) Specifies the name or ID of the mapping to stop.

Description

This command stops a single mapping. If the copy process is stopped, the target disk becomes unusable unless it already contained a complete image of the source (that is, unless the map had a progress of 100 as shown by the **lsfcmap** command). Before you can use the target disk, the mapping must once again be prepared and then reprocessed (unless the target disk already contained a complete image).

Only stand-alone mappings can be stopped using the **stopfcmap** command. Mappings that belong to a consistency group must be stopped using the **stopfcconsistgrp** command.

If the mapping is in the `idle_or_copied` state, the **stopfcmap** command has no effect and the mapping stays in the `idle_or_copied` state.

Note: Before SAN Volume Controller 4.2.0, the **stopfcmap** command always changed the mapping state to stopped and took the target VDisk offline. This change can break scripts that depend on the previous behavior.

The **split** option can be used when the mapping has progress of 100. It removes the dependency of any other mappings on the source VDisk. It might be used prior to starting another FlashCopy mapping whose target disk is the source disk of the mapping being stopped. Once the mapping has been stopped with the **split** option, the other mapping could then be started without the **restore** option.

An invocation example

```
svctask stopfcmap testmapone
```

The resulting output

No feedback

Chapter 15. Metro Mirror and Global Mirror commands

The following Copy Service commands enable you to work with the Metro Mirror and Global Mirror services that the SAN Volume Controller provides.

chpartnership

The **chpartnership** command modifies the bandwidth of the partnership between the local cluster and the remote cluster that is specified in the command. This affects the bandwidth that is available for background copy in a cluster partnership by either Metro Mirror or Global Mirror operations. This command can also be used to disable and re-enable the partnership, to permit the local cluster to be disconnected and then reconnected to the remote cluster.

Syntax

```
svctask -- chpartnership -- [-bandwidth bandwidth_in_mbps] [-start | -stop]
                             |
                             |
remote_cluster_id |
remote_cluster_name
```

Parameters

-bandwidth *bandwidth_in_mbps*

(Optional) Specifies the new bandwidth in megabytes per second (MBps). This bandwidth is used to cap the background remote copy progress. Set the bandwidth to the maximum rate that the remote copies should resynchronize at. Write operations from the host add to the use of the cluster link. If this parameter is set to a value that is greater than the intercluster links can sustain, the actual copy rate defaults to what is available on the link.

-start | -stop

(Optional) Starts or stops a Metro Mirror or Global Mirror partnership. To start or stop a partnership, run the **svctask chpartnership** command from either cluster.

remote_cluster_id | *remote_cluster_name*

(Required) Specifies the cluster ID or name of the remote cluster. The intracluster bandwidth cannot be modified, so if you enter the local cluster name or ID, an error occurs.

Description

This command modifies the bandwidth of the partnership between the local cluster and the remote cluster that is specified in the command. This affects the bandwidth that is available for a background copy in Metro Mirror or Global Mirror relationships, in the direction from the local to the remote cluster. To modify the background copy bandwidth in the other direction (remote cluster → local cluster), it is necessary to issue the corresponding **chpartnership** command to the remote cluster.

When you stop the cluster partnership, you temporarily disable the partnership and disconnect the local cluster from the remote cluster. The configuration is retained. The cluster partnership must be in either the `partially_configured_stopped` or `fully_configured_stopped` states to be started.

An invocation example

```
svctask chpartnership -bandwidth 20 cluster1
svctask chpartnership -stop cluster1
```

The resulting output

No feedback

chrconsistgrp

The **chrconsistgrp** command modifies the name of an existing Metro Mirror or Global Mirror consistency group.

Syntax

```
svctask -- chrconsistgrp -- -name -- new_name_arg -- [ rc_consist_group_name | rc_consist_group_id ]
```

Parameters

-name *new_name_arg*

(Required) Specifies the new name to assign to the consistency group.

rc_consist_group_name | *rc_consist_group_id*

(Required) Specifies the ID or existing name of the consistency group that you want to modify.

Description

This command changes the name of the specified consistency group.

An invocation example

Change the name of the consistency group called `rc_testgrp` to `rctestone`.

```
svctask chrconsistgrp -name rctestone rc_testgrp
```

The resulting output

No feedback

chrrelationship

The **chrrelationship** command enables you to modify certain attributes of an existing relationship, such as to add a relationship to a consistency group, to remove a relationship from a consistency group, and to change the name of the relationship. You can only change one attribute at a time per command submittal.

Syntax

```
svctask -- chrrelationship -- [ -name -- new_name_arg | -force | -consistgrp [ consist_group_id | consist_group_name ] ]
[ rc_rel_id | rc_rel_name ]
```

Parameters

-name *new_name_arg*

(Optional) Specifies a new label to assign to the relationship.

This parameter is required if you do not specify the **-consistgrp** or **-force** parameter.

-consistgrp *consist_group_id* | *consist_group_name*

(Optional) Specifies a new consistency group to assign the relationship to. Only relationships of the same copy type (Metro Mirror or Global Mirror) can be assigned to the same consistency group. You cannot use this parameter with the **-name**, or **-force** parameters.

This parameter is required if you do not specify the **-name** or **-force** parameter.

-force

(Optional) Specifies that you want the system to remove the relationship from a consistency group making the relationship a stand-alone relationship. You cannot use this parameter with the **-name** or **-consistgrp** parameters.

This parameter is required if you do not specify the **-name** or **-consistgrp** parameter.

rc_rel_name | *rc_rel_id*

(Required) Specifies the ID or name of the relationship.

Description

This command modifies the specified attributes of the supplied relationship, one attribute at a time; you cannot use more than one of the three optional parameters in the same command. In addition to changing the name of a consistency group, this command can be used for the following purposes.

- You can add a stand-alone relationship to a consistency group by specifying the **-consistgrp** parameter and the name or ID of the consistency group. The relationship and consistency group must be connected when the command is issued and must share the following components:
 - Master cluster
 - Auxiliary cluster
 - State (unless the group is empty)
 - Primary (unless the group is empty)
 - Type (unless the group is empty)

When the first relationship is added to an empty group, the group takes on the same state, primary (copy direction), and type (Metro or Global Mirror) as the relationship. Subsequent relationships must have the same state, copy direction, and type as the group in order to be added to it. A relationship can only belong to one consistency group.

- You can remove a relationship from a consistency group by specifying the **-force** parameter and the name or ID of the relationship. Although you do not have to specify or confirm the name of the consistency group, verify which group the relationship belongs to before you issue this command.

This form of the modify relationship command succeeds in the connected or disconnected states. If the clusters are disconnected the relationship is only removed from the consistency group on the local cluster, at the time the command is issued. When the clusters are reconnected the relationship is automatically removed from the consistency group on the other cluster. Alternatively, you can issue an explicit modify (**chrrelationship**) command to remove the relationship from the group on the other cluster while it is still disconnected.

Note: If you remove all relationships from the group, the relationship type is reset to **empty_group**. When you add a relationship to the empty group, the group again takes on the same type as the relationship.

- To move a relationship between two consistency groups, you must issue the **chrrelationship** command twice. Use the **-force** parameter to remove the relationship from its current group, and then use the **-consistgrp** parameter with the name of the new consistency group.

An invocation example

Change the name of the relationship `rccopy1` to `testrel`.

```
svctask chrrelationship -name testrel rccopy1
```

The resulting output

No feedback

An invocation example

Add relationship `rccopy2` to group called `newgroup`.
`svctask chrcrelationship -consistgrp newgroup rccopy2`

The resulting output

No feedback

An invocation example

Remove relationship `rccopy3` from whichever consistency group it is a member of.
`svctask chrcrelationship -force rccopy3`

The resulting output

No feedback

mkpartnership

The **mkpartnership** command establishes a one-way Metro Mirror or Global Mirror relationship between the local cluster and a remote cluster.

Syntax

To establish a fully functional Metro Mirror or Global Mirror partnership, you must issue this command to both clusters. This step is a prerequisite to creating Metro Mirror or Global Mirror relationships between VDisks on the clusters.

```
▶▶ svctask — — mkpartnership — — -bandwidth — bandwidth_in_mbps — —————▶
```



```
▶ remote_cluster_id —————▶
```

```
└─ remote_cluster_name ─┘
```

Parameters

-bandwidth *bandwidth_in_mbps*

(Required) Specifies the bandwidth, in megabytes per second (MBps), that is used by the background copy process between the clusters. It adjusts the bandwidth that is used by Metro Mirror or Global Mirror for the initial background copy process. Set the bandwidth to a value that is less than or equal to the bandwidth that can be sustained by the intercluster link. If the **-bandwidth** parameter is set to a higher value than the link can sustain, the background copy process uses the actual available bandwidth.

remote_cluster_id | *remote_cluster_name*

(Required) Specifies the cluster ID or name of the remote cluster. Issue the **svcinfolscandidate** command to list the remote clusters that are available. If two or more remote clusters have the same name and the name is included in this command, the command fails and it requests the ID of the cluster instead of the name.

Description

This command creates a one-way partnership between the local cluster and the remote cluster that you specify in the command. To create a two-way partnership, the equivalent **svctask mkpartnership** command must be issued from the other cluster. The maximum supported number of clusters in a partnership set is four.

Intercluster Mirror relationships can be created between primary VDisks in the local cluster and auxiliary VDisks in the remote cluster. Intracluster relationships can be created between VDisks that reside in a local cluster. The VDisks must belong to the same I/O group within the cluster.

Note: Background copy bandwidth can impact foreground I/O latency. To set the background copy bandwidth optimally, you must consider all three resources: the primary storage, intercluster link bandwidth, and secondary storage. Provision the most restrictive of these three resources between the background copy bandwidth and the peak foreground I/O workload.

An invocation example

```
svctask mkpartnership -bandwidth 20 cluster1
```

The resulting output

No feedback

mkrcconsistgrp

The **mkrcconsistgrp** command creates a new, empty Metro Mirror or Global Mirror consistency group. If the **-cluster** parameter is not specified, the consistency group is created on the local cluster only.

Syntax

```
svctask -- mkrcconsistgrp [ -name new_name ] [ -cluster cluster_id | cluster_name ]
```

Parameters

-name *new_name*

(Optional) Specifies a name for the new consistency group.

-cluster *cluster_id* | *cluster_name*

(Optional) Specifies the name or ID of the remote cluster. If **-cluster** is not specified, a consistency group is created only on the local cluster.

Description

This command creates a new consistency group. The ID of the new group is displayed after the command processes. The name must be unique across all consistency groups that are known to the clusters within this consistency group. If the consistency group involves two clusters, the clusters must be in communication throughout the create process.

The new consistency group does not contain any relationships and will be in the empty state. You can add Metro Mirror or Global Mirror relationships to the group using the **svctask chrcrelationship** command.

An invocation example

```
svctask mkrconsistgrp -name rc_testgrp
```

The resulting output

```
RC Consistency Group, id [255], successfully created
```

mkrelationship

The **mkrelationship** command creates a new Metro Mirror or Global Mirror relationship with virtual disks (VDisks) in the same cluster (intracluster relationship) or in two different clusters (intercluster relationship).

Syntax

```
svctask -- mkrelationship -- -master [ master_vdisk_id | master_vdisk_name ]
-- -aux [ aux_vdisk_id | aux_vdisk_name ] -- -cluster [ cluster_id | cluster_name ] -- -name [ new_name_id ]
-- -consistgrp [ consist_group_id | consist_group_name ] -- -sync -- -global
```

Parameters

-master *master_vdisk_id* | *master_vdisk_name*

(Required) Specifies the ID or name of the master virtual disk.

-aux *aux_vdisk_id* | *aux_vdisk_name*

(Required) Specifies the ID or name of the auxiliary virtual disk.

-cluster *cluster_id* | *cluster_name*

(Required) Specifies the ID or name of the remote cluster.

If you are creating an intracluster relationship, enter the ID of the local cluster. The VDisks in the relationship must belong to the same I/O group within the cluster.

If you are creating an intercluster relationship, enter the ID of the remote cluster. To create a relationship in two different clusters, the clusters must be connected at the time that the **svctask mkrelationship** command is received.

-name *new_name_id*

(Optional) Specifies a label to assign to the relationship.

-consistgrp *consist_group_id* | *consist_group_name*

(Optional) Specifies a consistency group that this relationship joins. If you do not supply the **-consistgrp** parameter, the relationship is created as a stand-alone relationship that can be started, stopped, and switched on its own.

Note: Metro and Global Mirror relationships cannot belong to the same consistency group. When the first relationship is added to the consistency group, the group takes on the same type as the relationship. Subsequently, only relationships of that type can be added to the consistency group.

-sync

(Optional) Specifies that you want the system to create a synchronized relationship. The **-sync** parameter guarantees that the master and auxiliary disks contain identical data at the point that the

relationship is created. You must ensure that the auxiliary disk is created to match the master disk and that no input transactions take place to either disk before you issue the create command. The initial background synchronization is skipped.

-global

(Optional) Specifies that you want the system to create a new Global Mirror relationship. If you do not specify the **-global** parameter, a Metro Mirror relationship is created instead.

Description

This command creates a new Metro Mirror or Global Mirror relationship. A Metro Mirror relationship defines the relationship between two virtual disks (VDisks): a master VDisk and an auxiliary VDisk. This relationship persists until it is deleted. The auxiliary virtual disk must be identical in size to the master virtual disk or the command fails, and if both VDIs are in the same cluster, they must both be in the same I/O group. The master and auxiliary cannot be in an existing relationship. Neither disk can be the target of a FlashCopy mapping. The command also returns the new relationship ID.

Metro Mirror relationships use one of the following copy types:

- A Metro Mirror copy ensures that updates are committed to both the primary and secondary VDIs before sending confirmation of I/O completion to the host application. This ensures that the secondary VDisk is synchronized with the primary VDisk in the event that a failover operation is performed.
- A Global Mirror copy allows the host application to receive confirmation of I/O completion before the updates are committed to the secondary VDisk. If a failover operation is performed, the host application must recover and apply any updates that were not committed to the secondary VDisk.

You can optionally give the relationship a name. The name must be a unique relationship name across both clusters.

The relationship can optionally be assigned to a consistency group. A consistency group ensures that a number of relationships are managed so that, in the event of a disconnection of the relationships, the data in all relationships within the group is in a consistent state. This can be important in, for example, a database application where data files and log files are stored on separate VDIs and consequently are managed by separate relationships. In the event of a disaster, the primary and secondary sites might become disconnected. As the disconnection occurs and the relationships stop copying data from the primary to the secondary site, there is no assurance that updates to the two separate secondary VDIs will stop in a consistent manner if the relationships that are associated with the VDIs are not in a consistency group.

For proper database operation, it is important that updates to the log files and the database data are made in a consistent and orderly fashion. It is crucial in this example that the log file VDisk and the data VDisk at the secondary site are in a consistent state. This can be achieved by putting the relationships that are associated with these VDIs into a consistency group. Both Metro Mirror and Global Mirror processing ensure that updates to both VDIs at the secondary site are stopped, leaving a consistent image based on the updates that occurred at the primary site.

If you specify a consistency group, both the group and the relationship must have been created using the same master cluster and the same auxiliary cluster. The relationship must not be a part of another consistency group. If the consistency group is *empty*, it acquires the type of the first relationship that is added to it. Therefore, each subsequent relationship that you add to the consistency group must have the same type.

If the consistency group is *not empty*, the consistency group and the relationship must be in the same state. If the consistency group is *empty*, it acquires the state of the first relationship that is added to it. If the state has an assigned copy direction, the direction of the consistency group and the relationship must match that direction.

If you do not specify a consistency group, a stand-alone relationship is created.

If you specify the **-sync** parameter, the master and auxiliary virtual disks contain identical data at the point when the relationship is created. You must ensure that the auxiliary is created to match the master and that no data movement occurs to either virtual disk before you issue the **svctask mkrrelationship** command.

If you specify the **-global** parameter, a Global Mirror relationship is created. Otherwise, a Metro Mirror relationship is created instead.

An invocation example

```
svctask mkrrelationship -master vdisk1 -aux vdisk2 -name rccopy1  
-cluster 0000020063432AFD
```

The resulting output

```
RC Relationship, id [28], successfully created
```

rmpartnership

The **rmpartnership** command removes a Metro Mirror or Global Mirror partnership on one cluster. Because the partnership exists on both clusters, it is necessary to run this command on both clusters to remove both sides of the partnership. If the command is run on only one cluster, the partnership enters a partially configured state on the other cluster.

Syntax

```
►► svctask — — rmpartnership — — [ remote_cluster_id | remote_cluster_name ] —►
```

Parameters

remote_cluster_id | *remote_cluster_name*

(Required) Specifies the cluster ID or the name of the remote cluster.

Description

This command deletes one half of a partnership on a cluster. To remove the entire partnership, you must run the command twice, once on each cluster.

Attention: Before running the **svctask rmpartnership** command, you must remove all relationships and groups that are defined between the two clusters. To display cluster relationships and groups, run the **svcinfo lsrelationship** and **svcinfo lsrconsistgrp** commands. To remove the relationships and groups that are defined between the two clusters, run the **svctask rmrrelationship** and **svctask rmrconsistgrp** commands.

An invocation example

```
svctask rmpartnership cluster1
```

The resulting output

```
No feedback
```

rmrconsistgrp

The **rmrconsistgrp** command deletes an existing Metro Mirror or Global Mirror consistency group.

Syntax

```
svctask -- rmrconsistgrp [ -force ] [ rc_consist_group_id | rc_consist_group_name ]
```

Parameters

-force

(Optional) Specifies that you want the system to remove all relationships belonging to a group before the consistency group is deleted. The relationships themselves are not deleted; they become stand-alone relationships.

Note: The **-force** parameter must be used to delete a consistency group when the consistency group has any Metro Mirror or Global Mirror relationships that is associated with it. If you do not use the **-force** parameter, the command fails.

rc_consist_group_id | *rc_consist_group_name*

(Required) Specifies the ID or the name of the consistency group to delete.

Description

This command deletes the specified consistency group. You can issue this command for any existing consistency group. If the consistency group is disconnected at the time that the command is issued, the consistency group is only deleted on the cluster that is connected. When the clusters reconnect, the consistency group is automatically deleted on the other cluster. Alternatively, if the clusters are disconnected, and you still want to remove the consistency group on both clusters, you can issue the **svctask rmrconsistgrp** command separately on both of the clusters.

If the consistency group is not empty, the **-force** parameter is required to delete the group. This removes the relationships from the consistency group before the group is deleted. These relationships become stand-alone relationships. The state of these relationships is not changed by the action of removing them from the consistency group.

An invocation example

```
svctask rmrconsistgrp rctestone
```

The resulting output

No feedback

rmrrelationship

The **rmrrelationship** command deletes an existing Metro Mirror or Global Mirror relationship.

Syntax

```
svctask -- rmrrelationship [ rc_rel_id | rc_rel_name ]
```

Parameters

rc_rel_id | *rc_rel_name*

(Required) Specifies the ID or the name of the relationship.

Description

This command deletes the relationship that is specified.

Deleting a relationship only deletes the logical relationship between the two virtual disks; it does not affect the virtual disks themselves.

If the relationship is disconnected at the time that the command is issued, the relationship is only deleted on the cluster where the command is being run. When the clusters reconnect, the relationship is automatically deleted on the other cluster. Alternatively, if the clusters are disconnected and if you still want to remove the relationship on both clusters, you can issue the **svctask rmrrelationship** command independently on both of the clusters.

If you delete an inconsistent relationship, the secondary virtual disk becomes accessible even though it is still inconsistent. This is the one case in which Metro or Global Mirror does not inhibit access to inconsistent data.

An invocation example

```
svctask rmrrelationship rccopy1
```

The resulting output

No feedback

starttrconsistgrp

The **starttrconsistgrp** command starts the Metro Mirror or Global Mirror consistency group copy process, sets the direction of copy if it is undefined, and optionally marks the secondary VDisks of the consistency group as clean.

Syntax

```
svctask -- starttrconsistgrp [-primary master | aux] [-force]
                                [-clean] rc_consist_group_id rc_consist_group_name
```

Parameters

-primary *master* | *aux*

(Optional) Specifies the copy direction by defining whether the master or auxiliary disk becomes the primary (source). This parameter is required when the primary is undefined if, for example, the consistency group is in the **Idling** state.

-force

(Optional) Specifies that you want the system to process the copy operation even if it might lead to a temporary loss of consistency while synchronization occurs. This parameter is required if the consistency group is in the **ConsistentStopped** state, but is not synchronized or is in the **Idling** state, but is not synchronized.

-clean

(Optional) Specifies that the VDisk that is to become a secondary is clean for each of the relationships belonging to the group; any changes made on the secondary VDisk are ignored, and only changes

made on the clean primary VDisk are considered during synchronization of the primary and secondary disks. The consistency group must be in an Idling (connected) state for this parameter to work.

Attention: This flag should only be used when the primary and secondary VDisks contain identical data. Otherwise, relationships that are not consistent are reported as consistent. Once this has been done there is no method to determine whether these VDisks ever reach a true consistent state until a full background copy can be carried out again.

rc_consist_group_id | *rc_consist_group_name*

(Required) Specifies the ID or name of the consistency group to start.

Description

This command starts a Metro Mirror or Global Mirror stand-alone consistency group.

This command can only be issued to a consistency group that is connected. For a consistency group that is idling, this command assigns a copy direction (primary and secondary roles) and begins the copy process. Otherwise, this command restarts a previous copy process that was stopped either by a stop command or by an I/O error.

If the resumption of the copy process leads to a period of time when the relationship is not consistent, then you must specify the **-force** parameter when you restart the relationship. This situation can arise if the relationship had been stopped and then further input transactions had been performed on the original primary disk of the relationship. When you use the **-force** parameter in this situation, the data on the secondary disk is not usable (because it is inconsistent) in a disaster recovery circumstance.

In the idling state, you must provide the **-primary** parameter. In other connected states, you can provide the **-primary** parameter, but it must match the existing setting.

The **-force** parameter is required if consistency would be lost by starting a copy operation. This can occur if write operations on either primary or secondary VDisks have taken place since the **ConsistentStopped** or **idling** state occurred. If the command is issued without the **-force** parameter in such circumstances, the command fails. In general, the **-force** parameter is required if the group is in one of the following states:

- Consistent_Stopped but not synchronized (sync=out_of_sync)
- Idling but not synchronized

The **-force** parameter is not required if the group is in one of the following states:

- Inconsistent_Stopped
- Inconsistent_Copying
- Consistent_Synchronized

However, the command does not fail if you specify the **-force** parameter.

The **-clean** parameter is used when a Metro Mirror or Global Mirror group is started and the secondary VDisks in this group are assumed to be clean. Clean in this sense, means that any changes that have been made at the secondary are ignored and only changes made at the primary are considered when synchronizing the primary and secondary VDisks. The **-clean** parameter can be used in the following scenario:

1. A consistency group is created with the **-sync** parameter. At this point, it does not matter if the primary and secondary contain the same data, even though the use of the **-sync** parameter implies that this is true.
2. A **stoprconsistgrp** command is issued with the **-access** parameter. This permits access to the secondary disk. Change recording begins at the primary.

3. An image of the primary disk is copied and loaded on to the secondary disk. It is permissible to allow updates to the primary disk during the image copy as this image can be only a fuzzy image of the primary disk.
4. A `starttrconsistgrp` command that specifies the `-primary master`, `-force`, and `-clean` parameters is issued. The auxiliary disk is marked as clean and changes on the master disk that have occurred since the relationship was stopped are copied to the auxiliary disk.
5. Once the background copy has completed, relationships in the group become consistent and synchronized.

An invocation example

```
svctask starttrconsistgrp rccopy1
```

The resulting output

No feedback

starttrrelationship

The `starttrrelationship` command starts the Metro Mirror or Global Mirror relationship copy process, sets the direction of copy if undefined, and optionally, marks the secondary VDisk of the relationship as clean. The relationship must be a stand-alone relationship.

Syntax

```

▶▶ svctask — — starttrrelationship — [ -primary master aux ] [ -force ] —————▶
                                     [ -clean ] [ rc_rel_id rc_rel_name ] —————▶▶

```

Parameters

`-primary master | aux`

(Optional) Specifies the copy direction by defining whether the master or auxiliary disk becomes the primary (source). This parameter is required when the primary is undefined if, for example, the relationship is in the idling state.

`-force`

(Optional) Specifies that you want the system to process the copy operation even if it might lead to a temporary loss of consistency while synchronization occurs. This parameter is required if the relationship is in the **Consistentstopped** state, but is not synchronized or in the Idling state, but is not synchronized.

`-clean`

(Optional) Specifies that the VDisk that is to become a secondary is clean; any changes made on the secondary VDisk are ignored, and only changes made on the clean primary VDisk are considered when synchronizing the primary and secondary disks. The relationship must be in an Idling (connected) state for this parameter to work.

Attention: This flag should only be used when the primary and secondary VDisks contain identical data. Otherwise, relationships that are not consistent are reported as consistent. Once this has been done there is no method to determine whether these VDisks ever reach a true consistent state until a full background copy can be carried out again.

rc_rel_id | rc_rel_name

(Required) Specifies the ID or name of the relationship that you want to start in a stand-alone relationship.

Description

The **startcrelationship** command starts a stand-alone relationship. The command fails if it is used to start a relationship that is part of a consistency group.

This command can only be issued to a relationship that is connected. For a relationship that is idling, this command assigns a copy direction (primary and secondary roles) and begins the copy process. Otherwise, this command restarts a previous copy process that was stopped either by a stop command or by some I/O error.

If the resumption of the copy process leads to a period of time when the relationship is not consistent, you must specify the **-force** parameter when you restart the relationship. This situation can arise if the relationship had been stopped, and then further input transactions had occurred on the original primary of the relationship. Even though you use of the **-force** parameter in this situation, the data on the secondary is not be useful for disaster recovery purposes because the relationship is in an inconsistent state.

In the idling state, you must provide the **-primary** parameter. In other connected states, you can provide the **-primary** parameter, but it must match the existing setting.

The **-force** parameter is required if consistency would be lost by starting a copy operation. This can occur if input transactions have occurred on either the primary or secondary VDisks since the **ConsistentStopped** or **Idling** state occurred. If the **startcrelationship** command is issued without the **-force** parameter in such circumstances, the command fails. In general, the **-force** parameter is required if the relationship is in one of the following states:

- ConsistentStopped but not synchronized
- Idling but not synchronized

The **-force** parameter is not required if the relationship is in one of the following states:

- InconsistentStopped
- InconsistentCopying
- ConsistentSynchronized

However, the command will not fail if you do specify the **-force** parameter.

The **-clean** parameter is used when a Metro Mirror or Global Mirror relationship is started and the secondary VDisk in the relationship is assumed to be clean; any changes made on the secondary VDisk are ignored, and only changes made on the clean primary VDisk are considered when synchronizing the primary and secondary disks. The **-clean** parameter can be used in the following circumstance:

1. A relationship is created with the **-sync** parameter specified. (At this point it does not matter if the primary and secondary disks contain the same data, even though the use of the **-sync** parameter implies that this is true).
2. A **svctask stopcrelationship** command is issued with the **-access** parameter specified. This permits access to the secondary disk. Change recording begins at the primary disk.
3. An image of the primary disk is copied and loaded on to the secondary disk. It is permissible to allow updates to the primary disk during the image copy as this image need only be a *fuzzy* image of the primary disk.
4. A **svctask startcrelationship** command that specifies the **-primary master**, **-force**, and **-clean** parameters is issued. The auxiliary disk is marked as clean and changes on the master disk that have occurred since the relationship was stopped are copied to the auxiliary disk.

5. Once the background copy has completed, the relationship becomes consistent and synchronized.

An invocation example

```
svctask starttrrelationship rccopy1
```

The resulting output

No feedback

stoprconsistgrp

The **stoprconsistgrp** command stops the copy process for a Metro Mirror or Global Mirror consistency group. This command can also be used to enable write access to the secondary VDisks in the group if the group is in a consistent state.

Syntax

```
▶▶ svctask — — stoprconsistgrp — [ -access ] [ rc_consist_group_id | rc_consist_group_name ] ▶▶
```

Parameters

-access

(Optional) Allows write access to consistent secondary VDisks in the consistency group.

rc_consist_group_id | *rc_consist_group_name*

(Required) Specifies the ID or the name of the consistency group to stop all processing for.

Description

This command applies to a consistency group. You can issue this command to stop processing on a consistency group that is copying from primary VDisks to secondary VDisks.

If the consistency group is in an inconsistent state, all copy operations stop and do not resume until you issue the **svctask starttrconsistgrp** command. For a consistency group in the `consistent_synchronized` state, this command causes a consistency freeze.

When a consistency group is in a consistent state (for example, in the `consistent_stopped`, `consistent_synchronized`, or `consistent_disconnected` state) you can issue the **access** parameter with the **stoprconsistgrp** command to enable write access to the secondary virtual disks within that group.

Table 18 shows consistency group initial and final states:

Table 18. **stoprconsistgrp** consistency group states

Initial state	Final state	Notes®
<code>inconsistent_stopped</code>	<code>inconsistent_stopped</code>	If access is specified, the command is rejected.
<code>inconsistent_copying</code>	<code>inconsistent_stopped</code>	If access is specified, the command is rejected with no effect and the relationship remains in the <code>inconsistent_copying</code> state.
<code>consistent_stopped</code>	<code>consistent_stopped</code>	If access is specified, the final state is idling.
<code>consistent_synchronized</code>	<code>consistent_stopped</code>	If access is specified, the final state is idling. If access is not specified, the final state is <code>consistent_stopped</code> .

Table 18. **stoprcconsistgrp** consistency group states (continued)

Initial state	Final state	Notes®
idling	idling	Remains in <code>idling</code> state whether access is specified or not.
idling_disconnected	unchanged	If specified without access , the relationship/group remains in <code>idling_disconnected</code> state. If the clusters reconnect, the relationship/group is in either <code>inconsistent_stopped</code> or <code>consistent_stopped</code> state.
inconsistent_disconnected	inconsistent_stopped	The command is rejected, with or without the access flag.
consistent_disconnected	consistent_stopped	The command is rejected if specified without access . If specified with access , the relationship/group moves to <code>idling_disconnected</code> .

An invocation example

```
svctask stoprcconsistgrp rccopy1
```

The resulting output

No feedback

stoprrelationship

The **stoprrelationship** command stops the copy process for a Metro Mirror or Global Mirror stand-alone relationship. You can also use this command to enable write access to a consistent secondary VDisk.

Syntax

```

▶▶ svctask — — stoprrelationship — [ -access ] [ rc_rel_id | rc_rel_name ]

```

Parameters

-access

(Optional) Specifies that the system allow write access to a consistent secondary VDisk.

rc_rel_id | *rc_rel_name*

(Required) Specifies the ID or the name of the relationship to stop all processing for.

Description

The **stoprrelationship** command applies to a stand-alone relationship. The command is rejected if it is addressed to a relationship that is part of a consistency group. You can issue this command to stop a relationship that is copying from primary to secondary VDisks.

If the relationship is in an inconsistent state, any copy operation stops and does not resume until you issue a **svctask startrelationship** command. For a relationship in the `consistent_synchronized` state, this command causes a consistency freeze.

When a relationship is in a consistent state – in the `consistent_stopped`, `consistent_synchronized`, or `consistent_disconnected` state – you can use the `access` parameter to enable write access to the secondary virtual disk. Table 19 provides consistency group initial and final states.

Table 19. `stopprrelationship` consistency group states

Initial state	Final state	Notes
<code>inconsistent_stopped</code>	<code>inconsistent_stopped</code>	If <code>access</code> is specified, the command is rejected.
<code>inconsistent_copying</code>	<code>inconsistent_stopped</code>	If <code>access</code> is specified, the command is rejected with no effect and the relationship remains in the <code>inconsistent_copying</code> state.
<code>consistent_stopped</code>	<code>consistent_stopped</code>	If <code>access</code> is specified, the final state is <code>idling</code> .
<code>consistent_synchronized</code>	<code>consistent_stopped</code>	If <code>access</code> is specified, the final state is <code>idling</code> . If <code>access</code> is not specified, the final state is <code>consistent_stopped</code> .
<code>idling</code>	<code>idling</code>	Remains in <code>idling</code> state whether <code>access</code> is specified or not.
<code>idling_disconnected</code>	<code>unchanged</code>	If specified without <code>access</code> , the relationship/group remains in <code>idling_disconnected</code> state. If the clusters reconnect, the relationship/group is in either <code>inconsistent_stopped</code> or <code>consistent_stopped</code> state.
<code>inconsistent_disconnected</code>	<code>inconsistent_stopped</code>	The command is rejected, with or without the <code>access</code> flag.
<code>consistent_disconnected</code>	<code>consistent_stopped</code>	The command is rejected if specified without <code>access</code> . If specified with <code>access</code> , the relationship/group moves to <code>idling_disconnected</code> .

An invocation example

```
svctask stopprrelationship rccopy1
```

The resulting output

No feedback

switchrconsistgrp

The `switchrconsistgrp` command reverses the roles of the primary and secondary virtual disks (VDisks) in a Metro Mirror or Global Mirror consistency group when that consistency group is in a consistent state. All the relationships in the consistency group are affected by this change.

Syntax

```

▶▶ svctask — — switchrconsistgrp — — -primary ——— master ——— rc_consist_group_id ———▶
                                     |  aux  |  rc_consist_group_name |

```


Parameters

-primary master | aux

(Required) Specifies whether the master or auxiliary side of the relationships in the group will become the primary VDisks.

rc_consist_group_id | rc_consist_group_name

(Required) Specifies the ID or name of the consistency group to switch.

Description

This command applies to a consistency group. It is normally issued to reverse the roles of the primary and secondary virtual disks in a consistency group, perhaps as part of a failover process that is associated with a disaster recovery event. Write access to the former primary VDisks is lost and write access to the new primary VDisks is acquired. This command is successful when the consistency group is in a connected, consistent state, and when reversing the direction of the relationships would not lead to a loss of consistency, for example, when the consistency group is consistent and synchronized. The consistency group must be in one of the following states in order for the **switchrconsistgrp** command to process correctly:

- ConsistentSynchronized
- ConsistentStopped and Synchronized
- Idling and Synchronized

The consistency group moves to the ConsistentSynchronized state after the successful completion of this command. If you specify the **-primary** parameter and it is the same as the current primary, the command has no effect.

An invocation example

```
svctask switchrconsistgrp -primary aux rccopy2
```

The resulting output

No feedback

switchrrelationship

The **switchrrelationship** command reverses the roles of primary and secondary virtual disks in a stand-alone Metro Mirror or Global Mirror relationship when that relationship is in a consistent state.

Syntax

```
▶▶— svctask — — switchrrelationship — — -primary — master — rc_rel_id —▶▶  
                                          aux — rc_rel_name —
```

Parameters

-primary master | aux

(Required) Specifies whether the master disk or the auxiliary disk is to be the primary.

rc_rel_id | rc_rel_name

(Required) Specifies the ID or the name of the relationship to switch.

Description

The **switchrrelationship** command applies to a stand-alone relationship. It is rejected if it is used to try to switch a relationship that is part of a consistency group. It is normally issued to reverse the roles of the primary and secondary virtual disk in a relationship perhaps as part of a failover process during a

disaster recovery event. Write access to the old primary disk is lost and write access to the new primary disk is acquired. This command is successful when the relationship is in a connected, consistent state, and when reversing the direction of the relationship does not lead to a loss of consistency; that is, when the relationship is consistent and synchronized. The relationship must be in one of the following states in order for the **switchrcrelationship** command to process correctly:

- ConsistentSynchronized
- ConsistentStopped and Synchronized
- Idling and Synchronized

The relationship moves to the **ConsistentSynchronized** state after the successful completion of this command. If you specify the **-primary** parameter with the current primary, the command has no effect.

An invocation example

```
svctask switchrcrelationship -primary master rccopy2
```

The resulting output

No feedback

Note: Migration activity on a single managed disk is limited to a maximum of 4 concurrent operations. This limit does not take into account whether the managed disk is the source or the destination target. If more than four migrations are scheduled for a particular managed disk, further migration operations are queued pending the completion of one of the currently running migrations. If a migration operation is stopped for any reason, a queued migration task can be started. However, if a migration is suspended, the current migration continues to use resources and a pending migration is not started. For example, the following setup is a possible initial configuration:

- MDiskGrp 1 has VDisk 1 created in it
- MDiskGrp 2 has VDisk 2 created in it
- MDiskGrp 3 has only one MDisk

With the previous configuration, the following migration operations are started:

- Migration 1 migrates VDisk 1 from MDiskGrp 1 to MDiskGrp 3, running with 4 threads.
- Migration 2 migrates VDisk 2 from MDiskGrp 2 to MDiskGrp 3, running with 4 threads.

Due to the previous limitations, the two migration operations do not always run at the same speed. MDiskGrp 3 has only one MDisk and the two migration operations have a total of 8 threads that are trying to access the one MDisk. Four threads are active. The remaining threads are in standby mode waiting to access the MDisk.

An invocation example

```
svctask migrateexts -vdisk vdisk4 -source mdisk4 -exts
64 -target mdisk6 -threads 4
```

The resulting output

No feedback

migratetoimage

The **migratetoimage** command migrates data from a VDisk (volume) (image mode or managed mode) onto a new image mode VDisk copy. The target disk does not have to be in the same MDisk group (storage pool) as the source disk.

Syntax

```
svctask -- migratetoimage -- -vdisk [source_vdisk_id] --
                                     [-copy id] [source_vdisk_name]
-- [-threads number_of_threads] -- -mdisk [unmanaged_target_mdisk_id]
                                     [unmanaged_target_mdisk_name]
| -- [-tier [generic_ssd] | generic_hdd] -- -mdiskgrp [managed_disk_group_id]
                                     [managed_disk_group_name]
```

Parameters

-vdisk *source_vdisk_id* | *name*

(Required) Specifies the name or ID of the source VDisk to be migrated.

-copy *id*

(Required if the specified VDisk has more than one copy) Specifies the VDisk copy to migrate from.

-threads *number_of_threads*

(Optional) Specifies the number of threads to use during the migration of extents. You can specify 1 - 4 threads. The default number of threads is 4.

-vdisk *vdisk_id* | *vdisk_name*

(Required) Specifies the virtual disk ID or name to migrate in to a new managed disk group.

Description

The **migratevdisk** command migrates the specified virtual disk into a new managed disk group; all the extents that make up the virtual disk are migrated onto free extents in the new managed disk group.

You can issue the **svcinfo lsmigrate** command to view the progress of the migration.

The process can be prioritized by specifying the number of threads to use during the migration. Using only one thread puts the least background load on the system.

The **migratevdisk** command fails if there are insufficient free extents on the targeted managed disk group for the duration of the command. To avoid this problem, do not issue new commands that use extents until the VDisk migration is completed.

The **migratevdisk** command fails if the target VDisk or source VDisk is offline. Correct the offline condition before attempting to migrate the VDisk.

An invocation example

```
svctask migratevdisk -vdisk 4 -mdiskgrp Group0 -threads 2
```

The resulting output

No feedback

Chapter 17. Tracing commands

Tracing commands capture information that can assist you with troubleshooting managed disks and virtual disks.

setdisktrace

Use the **setdisktrace** command to set a list of disks of a given type, to include in a disk trace.

Syntax

```
svctask -- setdisktrace -- -type [ mdisk | vdisk ] [ -set | -reset ]
[ -all | -objectid id_or_name_list ] [ -objectid id_or_name_list ]
```

Parameters

-type mdisk | vdisk

(Required) Specifies the object type for the disks.

-set

(Optional) Specifies the set argument. You cannot use the **-set** parameter with the **-reset** parameter.

-reset

(Optional) Specifies the reset argument. You cannot use the **-set** parameter with the **-reset** parameter.

-all

(Optional) Traces all disks of the specified type. You cannot use the **-all** parameter with the **-objectid** parameter.

-objectid id_or_name_list

(Optional) Specifies a list of one or more disk IDs or names. You cannot use the **-objectid** parameter with the **-all** parameter.

Description

The **setdisktrace** command marks the disks to be included in the next triggered trace.

The command is used with the **svctask settrace** command, which sets the options that result in a trace file and the data that is included in the trace file.

An invocation example

```
svctask setdisktrace -type mdisk -objectid
mdisk1:mdisk3:mdisk11:mdisk10:mdisk9:mdisk5 -reset
```

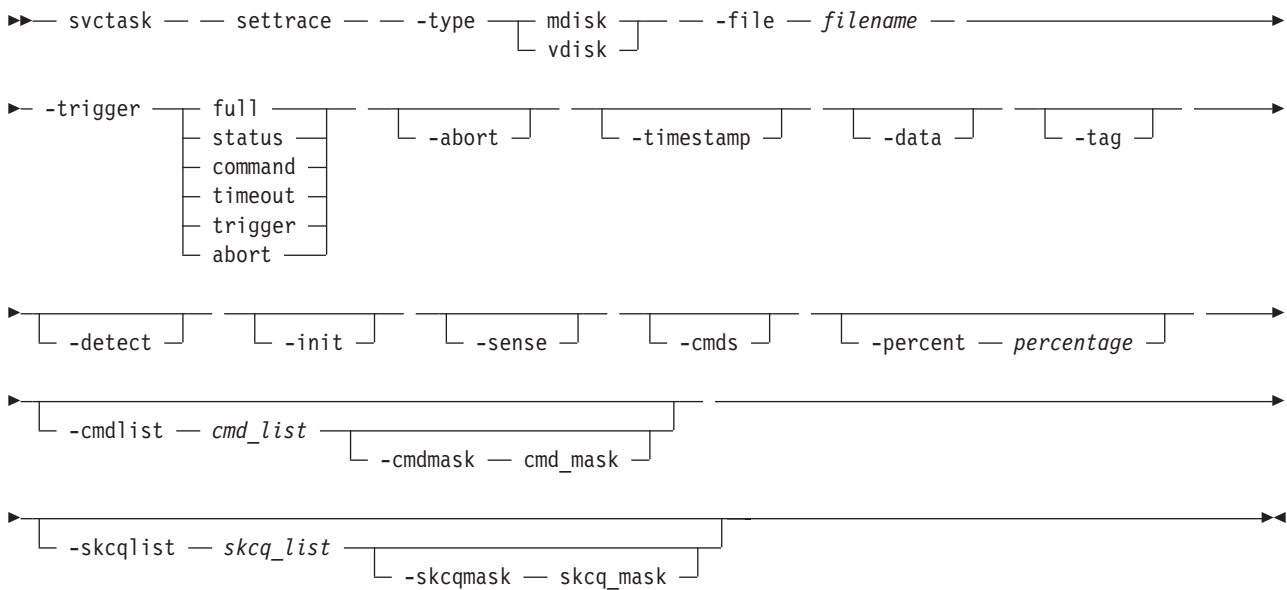
The resulting output

No feedback

settrace

The **settrace** command sets options to trace certain I/O operations through the system.

Syntax



Parameters

-type mdisk | vdisk

(Required) Specifies the type of objects to trace.

-file filename

(Required) Specifies the file name prefix for the trace file.

-trigger full | status | command | timeout | trigger | abort

(Required) Specifies an action for when the trace is started (triggered).

full Specifies to stop the trace when the trace buffer is full, for MDisks and VDIs.

status Sets a trigger for when the specified SCSI status (**-skcqlist**) is reported in sense data, for MDisks and VDIs.

command

Specifies a trigger for when the given SCSI command (**-cmdlist**) is sent, for MDisks and VDIs.

timeout

Sets a trigger for when a timeout occurs, for MDisks only.

trigger

Specifies to keep running until the trigger event, for MDisks only.

abort

Sets a trigger for when an abnormal end occurs, for VDIs only.

-abort

(Optional) Adds abnormal ending details to the trace, for VDIs only.

-timestamp

(Optional) Adds a time-stamp to each entry in the trace. A file name is created from the prefix plus a time-stamp. The file name is in the form *prefix_AAAAAA_YYMMDD_HHMMSS*, where *AAAAAA* is the panel name of the node generating the trace file.

-data

(Optional) Adds I/O data to the trace.

- tag**
(Optional) Adds CCB tags to the trace, for MDisks only.
- detect**
(Optional) Adds MDisk discovery details to the trace, for MDisks only.
- init**
(Optional) Adds MDisk initialization details to the trace, for MDisks only.
- sense**
(Optional) Adds SCSI sense data to the trace, for VDIsks only.
- cmds**
(Optional) Adds commands data to the trace, for VDIsks only.
- percent**
(Optional) Specifies the trigger point in the trace file, which determines the amount of data to collect after the trigger point. The default value is **50**, which places the trigger point in the middle of the trace file.
- cmdlist** *cmd_list*
(Optional) Adds the commands in the *cmd_list* to the trace file.
- cmdmask** *cmd_mask*
(Optional) Adds the commands in the *cmd_mask* to the trace file. The **-cmdmask** parameter must be used with the **-cmdlist** parameter.
- skcqlist** *skcq_list*
(Optional) Specifies an SKCQ list, which adds only those SKCQ details to the trace file.
- skcqmask** *skcq_mask*
(Optional) Specifies an SKCQ mask, which adds only those SKCQ details to the trace file. The **-skcqmask** parameter must be used with the **-skcqlist** parameter.

Description

The **settrace** command sets the various I/O tracing options for managed disks or virtual disks. When the relevant disk type trace is subsequently triggered, the options specify the data to be included in the trace file.

The file name specifies a file name prefix to use when you are generating a trace file. The system appends the node panel name and a timestamp to the file name.

A maximum of 10 trace files are kept on the cluster. When the eleventh trace is made, the oldest existing trace file is overwritten.

The directory can also hold files that are retrieved from other nodes. These files are not counted. The cluster deletes the oldest file to maintain the maximum number of files.

An invocation example

```
svctask settrace -type vdisk -file tracedump -trigger abort
  -percent 100 -abort -timestamp
```

The resulting output

No feedback

starttrace

Use the **starttrace** command to begin tracing I/O operations that are based on the option currently set for the specified object type and the list of disks to trace.

Syntax

```
svctask — — starttrace — — -type [ mdisk | vdisk ]
```

Parameters

-type mdisk | vdisk

Specifies the object type to trigger.

Description

This command starts the collection of I/O tracing information. The trace file is generated according to the options that you specified in the **svctask settrace** command. The disks that are traced are those that are identified in the list that is set by the **svctask setdisktrace** command.

The traces are written to the **/dumps/iotrace** directory. You can view the contents of this directory using the **svcinfo lsio tracedumps** command.

An invocation example

```
svctask starttrace -type vdisk
```

The resulting output

No feedback

stoptrace

Use the **stoptrace** command to stop tracing operations for the specified disk type.

Syntax

```
svctask — — stoptrace — — -type [ mdisk | vdisk ]
```

Parameters

-type mdisk | vdisk

(Required) Specifies the object type to stop tracing.

Description

This command stops the tracing of I/O operations for the specified object type. A trace file is not generated if the trigger options have not been met.

An invocation example

```
svctask stoptrace -type mdisk
```

The resulting output

No feedback

Chapter 18. Drive commands

Drive commands capture information that can assist you with managing drives.

applydrivesoftware

Use the `applydrivesoftware` command to upgrade drives.

Syntax

```
svctask -- applydrivesoftware -- -file name -- -type firmware
fpga
- -drive drive_id -- -force
```

Parameters

-file

Specifies the firmware upgrade file that exists in the `/home/admin/upgrade/` directory.

-type

Specifies the type of download.

-drive

Specifies a drive ID to apply software to, from the package provided.

-force

(Optional) This disables redundancy checking. In the unlikely event that a software installation causes the drive to fail, disabling redundancy checking might cause loss of data, or loss of access to data.

Note: This option is not recommended for redundant RAID configurations.

Description

This command upgrades drives. After you install a firmware package and run a single CLI, the system applies updates to the drive that has an update in that package. The system should stop if any problems occur.

An invocation example

```
svctask applydrivesoftware -file
drivemicrocodepackagev1 -type fpga -drive 4
```

chdrive

Use the `chdrive` command to change the drive properties.

Syntax

```
svctask -- chdrive -- -use unused | candidate | spare | failed -allowdegraded
- -task format | certify | recover -- drive_id
```

Parameters

-use

Describes the role of the drive:

- **unused**: the drive is not in use and will not be used as a spare
- **candidate**: the drive is available for use in an array
- **spare**: the drive can be used as a hot spare if required
- **failed**: the drive has failed.

Note: To create member drives, add the drives to arrays using the **charray** command.

-allowdegraded

(Optional) Permits permission for a change of drive to continue, even if a hotspare is not available.

-task

Causes the drive to perform a task:

- **format**: a drive is formatted for use in an array; only permitted when drive is a candidate or has failed validation
- **certify**: the disk is analyzed to verify the integrity of the data it contains; permitted for any drive that is a candidate, spare, or member
- **recover**: recover an offline SSD drive without losing data; permitted when the drive is offline because a build is required, or when the drive has failed validation

Note: You can track the drive progress using the **svcinfoldriveprogress** command.

drive_id

The identity of the drive.

Description

Use this command to change the drive properties.

```
svctask chdrive -use spare 1
```

lsdrive

Use the **lsdrive** command to display configuration information and drive VPD.

Syntax

```
▶▶— svcinfoldrive — --bytes— drive_id —▶▶
```

Parameters

drive_id

The identity of the drive.

Description

Use this command to display configuration information and drive VPD.

Note: Filtering should be permitted on all concise fields.

Table 20 on page 235 describes possible outputs.

Table 20. *lsdrive* output

Attribute	Value
id	The ID of the drive.
status	The summary status of the drive.
error_sequence_number	The error sequence number describing the cause of the drive status: <ul style="list-style-type: none"> • online: blank • degraded: populated if associated with an error • offline: must be populated <p>Note: Error sequence numbers indicate an entry in the event log. This includes entries that are both errors, and informational messages (for example, the drive is formatting).</p>
use	The current role of the drive: <ul style="list-style-type: none"> • unused: the drive which is not configured to be used by anything • candidate: the drive is available to be configured • spare: the drive is configured as a spare, to be used if the arrays fail members • member: the drive is configured as a member of an array • failed: the drive has been rejected, and is no longer available for use
UID	The unique ID reported by the drive.
tech_type	The drive technology used.
capacity	The capacity of disk, excluding quorum area.
block_size	The block size of the disk.
vendor_id	The manufacturer of the drive.
product_id	The product ID of the drive.
FRU_part_number	The FRU part number of the drive.
FRU_identity	The 11S number combining manufacturing part number and serial number.
RPM	The specified RPM of the disk.
firmware_level	Firmware level of the disk; blank if unknown.
FPGA_level	The FPGA level, if applicable. Blank if not applicable or unknown.
mdisk_id	The ID of the array MDisk that the drive is a member of.
mdisk_name	The name of the MDisk that the drive is a member of.
member_id	The ID of the MDisk array member.
enclosure_id	<ul style="list-style-type: none"> • If the drive is contained in an enclosure (not a node) and the slot position is known, this is the ID of the enclosure in which the drive is located. • If the drive is contained in a node (not an enclosure), this is blank. • If the enclosure ID has not been determined yet, this is blank.
slot_id	The slot_id of the drive in the enclosure or node. It can be referred to as the drive bay or location. This can be blank.
node_name	For a drive contained within a node, the node name where the drive is located. For a drive contained within an enclosure, blank.
node_id	For a drive contained within a node, the node ID where the drive is located. For a drive contained within an enclosure, blank.
quorum_id	The ID of quorum disk; blank if not quorum disk.
port_1_status	The connectivity status of the target for MDisk enumeration, with states.
port_2_status	

| **Concise invocation example:**

| `svcinfo lsdrive -delim :`

| **The resulting output:**

| id:status:error_sequence_number:use:tech_type:capacity:mdisk_id:mdisk_name:member_id:enclosure_id:slot_id
| 0:online::member:sas_ssd:20GB:0:mdisk0:0:1:2
| 1:offline:345:member:sas_ssd:20GB:0:mdisk0:0:1:3
| 2:online::member:sas_ssd:20GB:0:mdisk0:0:1:4

| **A detailed invocation example:**

| `svcinfo lsdrive 0`

| **The resulting output:**

| id:0
| status:online
| error_sequence_number:
| use:member
| UID:20000004cf4cd2c0
| tech_type:ssd
| capacity:20GB
| block_size:512
| vendor_id:IBM
| product_id:I8MR1337 W00Y4Y1
| FRU_part_number:AAAAAAA
| FRU_identity:11S1817115Y41337171001
| RPM:15000
| firmware_level:3.02
| FPGA_level:1.99
| mdisk_id:0
| mdisk_name:mdisk0
| member_id:0
| enclosure_id:1
| slot:2
| node_id:
| node_name:
| quorum_id:
| port_1_status:online
| port_2_status:online

| **Isdrivelba**

| Use the **Isdrivelba** command to map array MDisk logical block address (LBA) to a set of drives.

| **Syntax**

| `▶▶— svcinfo — lsdrivelba — -mdisklba — lba — -mdisk — mdisk_id — | mdisk_name —▶▶`

| **Parameters**

| **-mdisk**

| The ID of the MDisk.

| **-mdisklba**

| The logical block address (LBA) on the MDisk.

| **Description**

| This command maps the array MDisk logical block address (LBA) to a set of drives.

| Table 21 on page 237 describes possible outputs.

Table 21. *lsdrivelba* output

Attribute	Value
drive_id	The ID of drive; blank if no configured array member exists (for example, in a degraded array).
type	The type of information on the disk: <ul style="list-style-type: none"> • parity - LBA range contains parity (RAID levels 5 and 6 only) • qparity - LBA range contains qparity (RAID level 6 only) • data - LBA range contains data
drive_lba	The LBA on the drive.
drive_start	The start of range of LBAs (strip) on the drive.
drive_end	The end of range of LBAs (strip) on the drive.
mdisk_start	The start of range of LBAs (strip) on the array MDisk.
mdisk_end	The end of range of LBAs (strip) on the array MDisk.

This is an example of a five-member RAID-5 array with strip size of 256 KB:

An invocation example

```
svcinfolsdriovelba -mdisklba 0x000 -mdisk 2 -delim :
```

The resulting output

```
drive_id:type:drive_lba:drive_start:drive_end:mdisk_start:mdisk_end
0:data:0x0000000000000000:0x0000000000000000:0x0000000000000200:0x0000000000000000:0x0000000000000200
4:parity:0x0000000000000000:0x0000000000000000:0x0000000000000200:0x0000000000000000:0x0000000000000200
```

lsdriveprogress

Use the **lsdriveprogress** command to view the progress of various drive tasks.

Syntax

```
svcinfolsdriovelba --lsdriveprogress [drive_id]
```

Parameters

drive_id
 (Optional) The drive for which you want to view progress.

Description

The following outputs are possible:

drive_id
 The ID for the drive with the active task.

task The type of task:

- format
- certify
- rebuild

progress
 The percentage complete of the job.

| **estimated_completion_time**
| The estimated completion time (YYMMDDHHMMSS).

| **An invocation example**
| `svcinfo lsdriveprogress -delim :`

| **The resulting output**
| `drive_id:task:progress:estimated_completion_time`
| `0:format:10:091118131056`
| `9:certify:25:991231235959`

| **An invocation example**
| `svcinfo lsdriveprogress 9 -delim :`

| **The resulting output**
| `9:certify:25:991231235959`

| **triggerdrivedump**

| Use the **triggerdrivedump** command to collect support data from a disk drive. This data can help to understand problems with the drive, and does not contain any data that applications may have written to the drive.

| **Syntax**

| `▶▶— svctask — triggerdrivedump —drive_id—————▶◀`
|

| **Parameters**

| **drive_id**
| The ID of the drive to dump.

| **An invocation example**
| `svctask triggerdrivedump 1`

| **The resulting output**
| Dump file created on node id [2]

|

Chapter 19. Enclosure commands

Storwize® V7000 only: Enclosure commands capture information that can assist you with managing enclosures.

chenclosure

Use the **chenclosure** command to modify enclosure properties.

Syntax

```
| ▶▶ svctask — chenclosure —————→  
|                               | -identify —yes|no| | -managed —yes|no|  
|                               |  
| ▶ ————— enclosure_id —————→  
| | -id —enclosure_id|  
|
```

Parameters

Note: Optional parameters are mutually exclusive. Exactly one of the optional parameters must be set.

-identify *yes* | *no*

(Optional) Makes the identify LED start or stop flashing.

-managed *yes* | *no*

(Optional) Makes the enclosure into a managed or unmanaged enclosure.

-id *enclosure_id*

(Optional) Changes the enclosure ID after you replace the enclosure, and enables you to control what is on the front panel.

enclosure_id

The enclosure you want to modify.

Description

Use this command to modify enclosure properties.

To change the identity of enclosure 7 from 7 to 4:

```
svctask chenclosure -id 4 7
```

To change enclosure 1 to *unmanaged*:

```
svctask chenclosure -managed no 1
```

To make the identify LED on enclosure 1 stop flashing:

```
svctask chenclosure -identify no 1
```

chenclosurecanister

Use the **chenclosurecanister** command to modify the properties of an enclosure canister.

Syntax

```
svctask — chenclosurecanister — [ -excludesasport —yes|no— -port —1|2— -force —  
-identify —yes|no— ]  
-canister —canister_id—enclosure_id—
```

Note:

1. The **-port** and **-excludesasport** parameters must be specified together.
2. Exactly one of the optional parameters must be set.

Parameters

Note: Optional parameters are mutually exclusive.

-identify *yes|no*

(Optional) Change the state of fault LED to or from *slow_flashing*.

-excludesasport *yes|no*

(Optional) Excludes or includes the specified SAS port. You can use the **-force** flag if there are dependent VDIs.

Note: Using the **-force** flag might result in loss of access to your data.

-port *1|2*

(Optional) The SAS port to include or exclude.

canister_id

The canister you want to apply the change to.

enclosure_id

The enclosure that the canister is a member of.

Description

This command enables you to modify the properties of an enclosure canister.

To make the fault LED flash on canister 1 of enclosure 3:

```
svctask chenclosurecanister -identify yes -canister 1 3
```

To exclude SAS port 1 on canister 2 of enclosure 1:

```
svctask chenclosurecanister -excludesasport yes -port 1 -canister 2 1
```

chenclosureslot

Use the **chenclosureslot** command to modify the properties of an enclosure slot.

Syntax

```
svctask — chenclosureslot — [ -identify —yes|no—  
-exclude —yes|no— -port —port_id— -force — ]  
-slot —slot_id—enclosure_id—
```

| **Note:**

- | 1. Optional parameters are mutually exclusive.
- | 2. **-port** and **-force** may only be specified if **-exclude** is.
- | 3. Exactly one of the optional parameters must be set.
- | 4. The **-force** flag will only have an effect on the operation of **-exclude yes** .

| **Parameters**

| **-identify** *yes|no*

| Change the state of fault LED to or from *slow_flashing*.

| **-exclude** *yes|no*

| (Optional) Ensures that an enclosure slot port is excluded. The following list gives details of the options you can use with this parameter:

- | • **-exclude yes -port** *port_id* **-slot** *slot_id enclosure_id*: The port you specify with *port_id* will be excluded. If the current state of the port is *excluded_by_enclosure*, *excluded_by_drive*, or *excluded_by_cluster*, this command will appear to have no effect. However, if the current state of the port is *online*, then that state will change to *excluded_by_cluster*. The port will remain excluded until you rerun this command with *no* selected.

| **Attention:** This command will check for dependent volumes. If issuing this command would result in losing access to data, then the command will fail and an error message will display. You can use the **-force** flag to ignore these errors, but this could result in loss of access to data.

- | • **-exclude no -port** *port_id* **-slot** *slot_id enclosure_id* : The port will be put into *online* state, provided there are no other reasons to exclude the port. If you issue this command when the port is *online*, then it will have no effect. However, if you issue this command when the port is *excluded*, then the port state will do one of the following:
 - | – Change to *online* status immediately.
 - | – Change to *online* status after all other reasons for the port to be excluded have been removed.
- | • **-exclude yes|no -slot** *slot_id enclosure_id*: If you issue this command without defining a port, then the command will simultaneously act on both ports.

| **-port** *1|2*

| (Optional) The port on the canister to be excluded. If it is not specified, *-exclude* will act on both ports.

| **-slot** *slot_id*

| The slot ID.

| **-enclosure_id**

| The enclosure that the slot is a member of.

| **Description**

| This command enables you to modify the properties of an enclosure slot.

| Turn on the identify LED on slot 7 of enclosure 1:

```
| svctask chenclosureslot -identify yes -slot 7 1
```

| Force the exclusion of port 1 of slot 7 of enclosure 1:

```
| svctask chenclosureslot -exclude yes -port 1 -force -slot 7 1
```

| **Isenclosure**

| Use the **Isenclosure** command to view a summary of the enclosures.

Syntax

```
svcinfo -l enclosure enclosure_id -delim delimiter
```

Parameters

enclosure_id

Detailed information for the enclosure that you specify.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command enables you to view a summary of the enclosures (including current status information for canisters and power and cooling units, and other enclosure attributes). Table 22 shows the possible outputs:

Table 22. *lsenclosure* output

Attribute	Description
id	The ID of the enclosure.
status	Indicates if an enclosure is visible to the SAS network: <ul style="list-style-type: none">• online: a managed or unmanaged enclosure is visible• offline: a managed enclosure is not visible, and other fields hold their last known values.• degraded: if an enclosure is visible, but not down both strands
type	The type of enclosure: <ul style="list-style-type: none">• control• expansion
managed	Whether the enclosure is managed: <ul style="list-style-type: none">• yes• no
IO_group_id	The I/O group the enclosure belongs to; blank if canisters are connected to two different I/O groups.
IO_group_name	The I/O group the enclosure belongs to; blank if canisters are connected to two different I/O groups.
fault_LED	The status of the fault LED on the enclosure: <ul style="list-style-type: none">• on: a service action is required immediately on the enclosure or a component within the enclosure (including a canister, power unit, or non-spared drive).• slow_flashing: there is insufficient battery power to run I/O• off: there are faults on the enclosure or its components
identify_LED	The state of the identify LED <ul style="list-style-type: none">• off: the enclosure is not identified• slow_flashing: the enclosure is being identified

Table 22. *lsenclosure* output (continued)

Attribute	Description
error_sequence_number	Indicates the error log number of the highest priority error for this object. This is typically blank; however, if there is a problem (for example, the status has degraded), then it contains the sequence number of that error.
product_MTM	The product machine type and model.
serial_number	The serial number of the enclosure. This is the product serial number, which indicates the enclosure and its contents. The enclosure has its own serial number, which is embedded in the FRU_identity 11S data.
FRU_part_number	The FRU part number of the enclosure.
FRU_identity	The 11S serial number that combines the manufacturing part number and the serial number.
total_canisters	The maximum number of canisters for this enclosure type.
online_canisters	The number of canisters contained in this enclosure that are online.
total_PSUs	The number of power and cooling units in this enclosure.
online_PSUs	The number of PSUs contained in this enclosure that are online.
drive_slots	The number of drive slots in the enclosure.

An invocation example

```
svcinfolenclosure -delim :
```

The resulting output

```
id:status:type:managed:IO_group_id:IO_group_name:product_MTM:serial_number:total_canisters:online_canisters:
total_PSUs:online_PSUs:drive_slots1:degraded:expansion:no:0:io_grp0:2076-224:66G0083:2:2:2:24
```

lsenclosurebattery

Use the **lsenclosurebattery** command to display information about the batteries in the enclosure PSUs.

Syntax

```
svcinfolenclosurebattery -delim delimiter
```

```
-battery battery_id enclosure_id
```

Parameters

enclosure_id

(Optional) Lists the batteries for the specified enclosure.

-battery *battery_id*

(Optional) Provides a detailed view of the specified enclosure battery. Valid only when an enclosure is specified.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all

items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays information about the batteries in the enclosure PSUs. The concise view will show a line for each battery slot in every control enclosure, regardless of whether they exist. Batteries will not be shown for expansion enclosures. Table 23 shows possible outputs.

Table 23. *lsenclosurebattery* outputs

Attribute	Description
enclosure_id	The identity of the enclosure that contains the battery.
battery_id	Identifies the battery in the enclosure.
status	The status of the battery: <ul style="list-style-type: none"> • online: the battery is present and working as usual • degraded: the battery is present but not working as usual • offline: the battery could not be detected
charging_status	The charging state of the battery: <ul style="list-style-type: none"> • idle: the battery is not charging nor discharging • charging: the battery is charging • reconditioning: the battery is reconditioning itself, by being discharged and then recharged
recondition_needed	The battery needs to be reconditioned; however, this cannot be done because of one or more errors.
percent_charged	Indicates the charge of battery, in a percentage.
end_of_life_warning	The battery is reaching its end of life warning, and will need to be replaced: <ul style="list-style-type: none"> • yes • no
FRU_part_number	The FRU part number of the battery.
FRU_identity	The 11S number, combining the manufacturing part number and the serial number.
firmware_level	The battery firmware version.
error_sequence_number	Indicates the error log (or event log) number of the highest priority error for this object. This is typically blank; however, if there is a problem (for example, the status is degraded), then it contains the sequence number of that error event.

An invocation example

```
svcinfo lsenclosurebattery -delim :
```

The resulting output

```
enclosure_id:battery_id:status:charging_status:recondition_needed:percent_charged:end_of_life_warning
1:1:offline:idle:no:0:no
1:2:offline:idle:no:0:no
```

lsenclosurecanister

Use the **lsenclosurecanister** command to view a detailed status for each canister in an enclosure.

Syntax

```

➤➤ svcinfo — lsenclosurecanister — [ -canister — canister_id ] [ enclosure_id ]
|
| [ -delim — delimiter ]

```

Parameters

enclosure_id

Lists the canisters for the specified enclosure.

-canister *canister_id*

Valid only when the *enclosure_id* is specified. Provides a detailed view of the canister for the specified enclosure.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command enables you to view a detailed status for each canister in an enclosure. Table 24 shows the possible outputs:

Table 24. *lsenclosurecanister* output

Attribute	Description
<i>enclosure_id</i>	The identity of the enclosure that contains the canister.
<i>canister_id</i>	Identifies which of the canisters in the enclosure this is.
<i>status</i>	The status of the canister: <ul style="list-style-type: none"> • online: the canister is present and working normally. • degraded: the canister is present but not working normally • offline: the canister could not be detected.
<i>type</i>	The type of canister: <ul style="list-style-type: none"> • node • expansion
<i>node_id</i>	The node that corresponds to this canister; blank if the canister is not a node, or if the node is not part of the cluster.
<i>node_name</i>	The node that corresponds to this canister; blank if the canister is not a node, or if the node is not part of the cluster.
<i>FRU_part_number</i>	The FRU part number of the canister.
<i>FRU_identity</i>	The 11S number that combines the manufacturing part number and the serial number.
<i>WWNN</i>	The fibre channel WWNN of the canister (node canisters only).
<i>firmware_level</i>	The firmware level of the SCSI Enclosure Services (SES) code running on the canister.
<i>temperature</i>	(0 to 245) The temperature of the canister (in degrees Celsius). If the temperature goes below 0, 0 will be displayed.

Table 24. `lsenclosurecanister` output (continued)

Attribute	Description
fault_LED	The state of the combined fault and identify LEDs: <ul style="list-style-type: none"> off: no fault slow_flashing: identify mode <p>Note: When the LED is in identify mode, it conceals whether there is a fault present, because it always flashes. When you remove it from identity mode, the LED will become on or off.</p> <ul style="list-style-type: none"> on: fault
SES_status	The SCSI status of the connection between the device and the canister: <ul style="list-style-type: none"> online offline
error_sequence_number	Indicates the error log number of the highest priority error for this object. This is typically blank; however, if there is a problem (for example, the status is degraded), then it contains the sequence number of that error.
SAS_port_1_status	Indicates if there is damage to the cable between SAS ports: <ul style="list-style-type: none"> online offline excluded: are logged in, but cannot communicate with the canister degraded: the SAS cable not fully functional
SAS_port_2_status	Indicates if there is damage to the cable between SAS ports: <ul style="list-style-type: none"> online offline excluded: are logged in, but cannot communicate with the canister degraded: the SAS cable not fully functional

An invocation example

```
svcinfolenclosurecanister -delim :
```

The resulting output

```
enclosure_id:canister_id:status:type:node_id:node_name
1:1:degraded:expansion:1:node1
```

A detailed example

```
svcinfolenclosurecanister -canister 1 1
```

The resulting output

```
enclosure_id 1
canister_id 1
status online
type node
node_id 1
node_name node1
FRU_part_number AAAAAAA
FRU_identity 11S1234567Y12345678901
WWNN 5005076801005F94
firmware_level XXXXXXXXXX
temperature 23
fault_LED flashing
SES_status online
error_sequence_number
SAS_port_1_status online
SAS_port_2_status online
```


lsenclosurepsu

Use the **lsenclosurepsu** command to view information about each power-supply unit (PSU) in the enclosure.

Syntax

```
svcinfo lsenclosurepsu [-psu psu_id enclosure_id] [-delim delimiter]
```

Parameters

enclosure_id

(Optional) Lists the PSUs for the specified enclosure.

-psu *psu_id*

(Optional) Valid only when the **enclosure_id** is specified. Provides a detailed view of the PSU for the specified enclosure.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command enables you to view information about each power-supply unit (PSU) in the enclosure. Table 25 shows the possible outputs:

Table 25. *lsenclosurepsu* output

Attribute	Description
enclosure_id	The ID of the enclosure containing the PSU.
psu_id	The ID of the PSU in the enclosure.
status	The status of the power and cooling unit in the enclosure: <ul style="list-style-type: none">• online: a PSU is present and working normally• offline: a PSU cannot be detected• degraded: a PSU is present but not working normally
AC_failed	<ul style="list-style-type: none">• on: If the AC, DC, and fan LEDs are all on, then there is a PSU fault. If only the AC LED is on, then there is no AC power.• off: The AC power is OK.
DC_failed	<ul style="list-style-type: none">• on: If the AC, DC, and fan LEDs are all on, then there is a PSU fault. If only the DC LED is on, then there is no DC power.• off: The DC power is OK.
fan_failed	<ul style="list-style-type: none">• on: If the AC, DC, and fan LEDs are all on, then there is a PSU fault. If only the fan LED is on, then there is a fan failure.• off: The fans in this PSU are OK.

| Table 25. *lsenclosurepsu* output (continued)

Attribute	Description
redundant	Indicates if you can remove this power supply: <ul style="list-style-type: none"> • If the PSU is on an expansion enclosure, then the other PSU must be online. • If the PSU is on a control enclosure, then the other PSU must be online, and the battery on that PSU must contain enough charge to allow the canisters to dump state and cache data before shutting down.
error_sequence_number	Indicates the error log (or event log) number of the highest priority error for this object. This is typically blank; however, if there is a problem (for example, the status is degraded), then it contains the sequence number of that error event.
FRU_part_number	The FRU part number of the PSU.
FRU_identity	The 11S number, combining the manufacturing part number and the serial number.

| **An invocation example**

| `svcinfolenclosurepsu -delim :`

| **The resulting output**

| enclosure_id:PSU_id:status
| 1:1:degraded

lsenclosureslot

Use the **lsenclosureslot** command to view information about each drive slot in the enclosure.

Syntax

```

▶▶ svcinfolenclosureslot [-slot slot_id enclosure_id]
|
| [-delim delimiter]

```

Parameters

- | *enclosure_id*
Lists slots for that enclosure.
- | **-slot** *slot_id*
Valid only when an enclosure is specified. Gives detailed view for that enclosure slot.
- | **-delim** *delimiter*
(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command enables you to view information about each drive slot in the enclosure, such as whether a drive is present, and the port status for that drive. Table 26 on page 249 shows the possible outputs:

Table 26. `lsenclosureslot` output

Attribute	Description
<code>enclosure_id</code>	The identity of the enclosure which contains the drive slot.
<code>slot_id</code>	Identifies which of the drive slots in the enclosure this is.
<code>port_1_status</code>	The status of enclosure slot port 1. If the port is bypassed for multiple reasons, only one is shown. In order of priority, they are: <ul style="list-style-type: none"> • <code>online</code>: enclosure slot port 1 is online • <code>excluded_by_drive</code>: the drive excluded the port • <code>excluded_by_enclosure</code>: the enclosure excluded the port • <code>excluded_by_cluster</code>: the cluster has excluded the port
<code>port_2_status</code>	The status of enclosure slot port 2. If the port is bypassed for multiple reasons, only one is shown. In order of priority, they are: <ul style="list-style-type: none"> • <code>online</code>: enclosure slot port 2 is online • <code>excluded_by_drive</code>: the drive excluded the port • <code>excluded_by_enclosure</code>: the enclosure excluded the port • <code>excluded_by_cluster</code>: the cluster has excluded the port
<code>fault_LED</code>	The state of the combined fault and identify LEDs: <ul style="list-style-type: none"> • <code>off</code>: no fault • <code>slow_flashing</code>: identify mode <p>Note: When the LED is in identify mode, it conceals whether there is a fault present, because it always flashes. When you remove it from identify mode, the LED will become on or off.</p> <ul style="list-style-type: none"> • <code>on</code>: fault
<code>powered</code>	Indicates whether the slot is powered on. <ul style="list-style-type: none"> • <code>yes</code> • <code>no</code>
<code>drive_present</code>	Indicates if a drive is in the slot. The drive can be working, dead, or powered off. <ul style="list-style-type: none"> • <code>yes</code> (present) • <code>no</code> (empty)
<code>drive_id</code>	The ID of the drive in the slot; blank if there is no drive present, or if there is a drive present but it is offline and unmanaged.
<code>error_sequence_number</code>	Indicates the error log number of the highest priority error for this object. This is typically blank; however, if there is a problem (for example, the status is degraded), then it contains the sequence number of that error.

An invocation example

```
svcinfolenclosureslot -delim :
```

The resulting output

```
enclosure_id:slot_id:port_1_status:port_2_status:drive_present:drive_id
1:1:online:online:yes:22
1:2:online:online:yes:23
1:3:online:online:yes:19
1:4:online:online:yes:7
1:5:online:online:yes:10
1:6:online:online:yes:18
1:7:online:online:yes:20
1:8:online:online:yes:16
1:9:online:online:yes:12
1:10:online:online:yes:11
1:11:online:online:yes:21
```

```

| 1:12:online:online:yes:9
| 1:13:online:online:yes:14
| 1:14:online:online:yes:5
| 1:15:online:online:yes:15
| 1:16:online:online:yes:13
| 1:17:online:online:yes:6
| 1:18:online:online:yes:17
| 1:19:online:online:yes:4
| 1:20:online:online:yes:1
| 1:21:online:online:yes:8
| 1:22:online:online:yes:0
| 1:23:online:online:yes:3
| 1:24:online:online:yes:2

```

| triggerenclosedump

| Use the **triggerenclosedump** command to force the specified enclosure or enclosures to dump data.

| Syntax

```

| ▶▶— svctask — triggerenclosedump —————▶▶
|                                     |
|                                     | -port —port_id— -iogrp —iogrp_id_or_name—
|                                     | -enclosure —enclosure_id—
|                                     |

```

| Note:

- | 1. You can only use one of the optional parameters (**-port** or **-enclosure**).
- | 2. If **-port** is specified, **-iogrp** must also be specified.
- | 3. If **-iogrp** is specified, **-port** must also be specified.

| Parameters

- | **-port** *port_id*
 | If the system is wired correctly, this value is identical to the ID of the chain with the enclosures you want to dump. If the system is wired incorrectly, all the enclosures connected to port *port* of either node canister will be dumped.
- | **-iogrp** *iogrp_id_or_name*
 | The ID or name of the I/O group the control enclosure belongs to.
- | **-enclosure** *enclosure_id*
 | The ID of the enclosure you want to dump.

| Description

| This command requests the canisters in the enclosure or enclosures specified to dump data. The dumped data is subsequently collected and moved to /dumps/enclosure on the nodes that are connected to the enclosure. There will be one file for each canister successfully dumped and they may be located on different nodes. Dumps are for use by IBM support, which has the tools to interpret the dump data. Use the cpdumps command to copy the files from the system. This command does not disrupt access to the enclosures.

| **To trigger enclosure dumps from all enclosures connected to port 1 of the control enclosure in iogrp 2:**

```
| svctask triggerenclosedump -port 1 -iogrp 2
```

| **To trigger enclosure dumps from enclosure 5:**

```
| svctask triggerenclosedump -enclosure 5
```

Chapter 20. Attributes of the `-filtervalue` parameters

The `-filtervalue` parameter filters a view that is based on specific attribute values that relate to each object type. You can combine multiple filters to create specific searches, for example, `-filtervalue name=fred:status=online`. The help (`-filtervalue?`) specifies the attributes that are available for each object type.

The `-filtervalue` parameter must be specified with `attrib=value`. The `-filtervalue?` and `-filtervalue` parameters cannot be specified together.

Note: The qualifier characters left bracket (<) and right bracket (>) must be enclosed within double quotation marks (""). For example, `-filtervalue vdisk_count "<"4 or port_count ">"1`. It is also valid to include the entire expression within double quotation marks. For example, `-filtervalue "vdisk_count<4"`

When an attribute requires the `-unit` parameter, it is specified after the attribute. For example, `-filtervalue capacity=24 -unit mb`. The following input options are valid for the `-unit` parameter:

- **b** (bytes)
- **mb** (Megabytes)
- **gb** (Gigabytes)
- **tb** (Terabytes)
- **pb** (Petabytes)

Capacity values displayed in units other than bytes might be rounded. When filtering on capacity, use a unit of bytes, `-unit b`, for exact filtering.

Table 27 provides a list of valid filter attributes, as well as descriptions, qualifiers and wildcards for each object type.

You can use the asterisk (*) character as a wildcard character when names are used. The asterisk character can be used either at the beginning or the end of a text string, but not both. Only one asterisk character can be used in a `-filtervalue` parameter.

Table 27. Valid filter attributes

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
cluster	<i>cluster_name</i> or <i>name</i>	=	Yes	The cluster name.
	<i>cluster_unique_id</i> or <i>id</i>	=, <, <=, >, >=	No	The cluster ID.

Table 27. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
node	<i>node_name</i> or <i>name</i>	=	Yes	The node name.
	<i>id</i>	=, <, <=, >, >=	No	The node ID.
	<i>status</i>	=	No	The status of the node. The following values are valid for node <i>status</i> : <ul style="list-style-type: none"> • adding • deleting • online • offline • pending
	<i>IO_group_name</i>	=	Yes	The I/O group name.
	<i>IO_group_id</i>	=, <, <=, >, >=	No	The I/O group ID.
	<i>hardware</i>	=	No	The following values are valid for <i>hardware</i> type: 8F2 , 8F4 , 8G4 , CF8 , and 8A4 .
io_grp	<i>HWS_name</i> or <i>name</i>	=	Yes	The I/O group name.
	<i>HWS_unique_id</i> or <i>id</i>	=, <, <=, >, >=	No	The I/O group ID.
	<i>node_count</i>	=, <, <=, >, >=	No	The number of nodes in the I/O group.
	<i>host_count</i>	=, <, <=, >, >=	No	The number of hosts associated with the io_grp.
controller	<i>controller_id</i> or <i>id</i>	=, <, <=, >, >=	No	The controller ID.

Table 27. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
mdisk	<i>name</i>	=	Yes	The name of the MDisk.
	<i>id</i>	=, <, <=, >, >=	No	The ID of the MDisk.
	<i>controller_name</i>	=	Yes	The name of the controller the MDisk belongs to.
	<i>status</i>	=	No	The status of the MDisk. The following values are valid for MDisk <i>status</i> : <ul style="list-style-type: none"> • online • degraded_ports • degraded_paths • offline • excluded
	<i>mode</i>	=	No	The mode of the MDisk. The following values are valid for MDisk <i>mode</i> : <ul style="list-style-type: none"> • unmanaged • managed • image
	<i>mdisk_grp_name</i>	=	Yes	The MDisk group name.
	<i>mdisk_grp_id</i>	=, <, <=, >, >=	No	The MDisk group ID.
	<i>capacity</i>	=, <, <=, >, >=	No	The capacity. Requires the -unit parameter.
mdiskgrp	<i>tier</i>	=	No	The tier information being reported: <ul style="list-style-type: none"> • generic_hdd • generic_ssd
	<i>name</i>	=	Yes	The MDisk group name.
	<i>storage_pool_id</i> or <i>id</i>	=, <, <=, >, >=	No	The MDisk group ID.
	<i>mdisk_count</i>	=, <, <=, >, >=	No	The number of MDisks in the group.
	<i>vdisk_count</i>	=, <, <=, >, >=	No	The number of VDIs in the group.
	<i>status</i>	=	No	The status of the MDisk group. The valid input options are online , degraded_ports , degraded_paths , excluded , and offline .
	<i>extent_size</i>	=, <, <=, >, >=	No	The extent size. (MB)
	<i>easy_tier</i>	=	No	Determines if Easy Tier is permitted to manage the storage pool: <ul style="list-style-type: none"> • on • off
<i>easy_tier_status</i>	=	No	Determines if automatic data placement function on a storage pool is activated: <ul style="list-style-type: none"> • active • inactive 	

Table 27. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
vdisk	<i>vdisk_name</i> or <i>name</i>	=	Yes	The name of the VDisk.
	<i>vdisk_id</i> or <i>id</i>	=, <, <=, >, >=	No	The ID of the VDisk.
	<i>IO_group_name</i>	=	Yes	The name of the I/O group.
	<i>IO_group_id</i>	=, <, <=, >, >=	No	The ID of the I/O group.
	<i>status</i>	=	No	The status of the VDisk. The valid input options for VDisk status are online , degraded , and offline .
	<i>mdisk_grp_name</i>	=	Yes	The MDisk group name.
	<i>mdisk_grp_id</i>	=, <, <=, >, >=	No	The MDisk group ID.
	<i>capacity</i>	=, <, <=, >, >=	No	The capacity. Requires the -unit argument.
	<i>type</i>	=	No	The VDisk type. The valid value options are seq , striped , and image .
	<i>FC_name</i>	=	Yes	The FlashCopy mapping name.
	<i>FC_id</i>	=, <, <=, >, >=	No	The FlashCopy mapping ID.
	<i>fc_map_count</i>	=, <, <=, >, >=	No	The number of VDisk mappings (either source or target).
	<i>copy_count</i>	=, <, <=, >, >=	No	The number of VDisk mirrored copies.
<i>RC_name</i>	=	Yes	The Metro Mirror relationship name.	
<i>RC_id</i>	=, <, <=, >, >=	No	The Metro Mirror relationship ID.	

Table 27. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
vdisk_copy	<i>primary</i>	=	No	Indicates that this copy is the primary copy. The valid values are yes and no .
	<i>status</i>	=	No	The status of the MDisk group. Valid values are online , degraded , or offline .
	<i>sync</i>	=	No	Indicates whether the VDisk copy is synchronized. Valid values are true or false .
	<i>mdisk_grp_name</i>	=	Yes	The name of the MDisk group.
	<i>mdisk_grp_id</i>	=, <, <=, >, >=	No	The ID of the MDisk group.
	<i>type</i>	=	No	The type of the VDisk copy. The valid values are seq , striped , or image .
	<i>easy_tier</i>	=	No	Determines if Easy Tier is permitted to manage the storage pool: <ul style="list-style-type: none"> • on • off
	<i>easy_tier_status</i>	=	No	Determines if automatic data placement function on a storage pool is activated: <ul style="list-style-type: none"> • active • measured • inactive
se_vdiskcopy	<i>mdisk_grp_id</i>	=, <, <=, >, >=	No	The ID of the MDisk group.
	<i>mdisk_grp_name</i>	=	Yes	The name of the MDisk group.
	<i>overallocation</i>	=	No	The percentage of overallocation, which is displayed as a number.
	<i>autoexpand</i>	=	No	Autoexpand flags. The valid values are on and off .
	<i>grainsize</i>	=, <, <=, >, >=	No	Space-efficient grain size. The valid values are 32 , 64 , 128 , or 256 .
host	<i>host_name</i> or <i>name</i>	=	Yes	The host name.
	<i>host_id</i> or <i>id</i>	=, <, <=, >, >=	No	The host ID.
	<i>port_count</i>	=, <, <=, >, >=	No	The number of ports.
	<i>iogrp_count</i>	=, <, <=, >, >=	No	The number of I/O groups that are associated with the host.

Table 27. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
fcmap	<i>FC_mapping_name</i> or <i>name</i>	=	Yes	The FlashCopy mapping name.
	<i>FC_id</i> or <i>id</i>	=, <, <=, >, >=	No	The FlashCopy mapping ID.
	<i>source_vdisk_name</i>	=	Yes	The source VDisk name.
	<i>source_vdisk_id</i>	=, <, <=, >, >=	No	The source VDisk ID.
	<i>target_vdisk_name</i>	=	Yes	The target VDisk name.
	<i>target_vdisk_id</i>	=, <, <=, >, >=	No	The target VDisk ID.
	<i>group_name</i>	=	Yes	The consistency group name.
	<i>group_id</i>	=, <, <=, >, >=	No	The consistency group ID.
	<i>status</i>	=	No	The mapping status. The following values are valid for <i>fcmap status</i> : <ul style="list-style-type: none"> • idle_or_copied • preparing • prepared • copying • stopped • suspended • stopping • empty
<i>copy_rate</i>	=, <, <=, >, >=	No	The background copy rate.	
fcconsistgrp	<i>name</i>	=	Yes	The consistency group name.
	<i>FC_group_id</i> or <i>id</i>	=, <, <=, >, >=	No	The consistency group ID.
	<i>status</i>	=	No	The consistency group status. The following values are valid for <i>fcconsistgrp status</i> : <ul style="list-style-type: none"> • idle_or_copied • preparing • prepared • copying • stopped • suspended • stopping • empty

Table 27. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
rcrelationship	<i>RC_rel_id</i> or <i>id</i>	=, <, <=, >, >=	No	The Metro Mirror relationship ID.
	<i>RC_rel_name</i> or <i>name</i>	=	Yes	The Metro Mirror relationship name.
	<i>master_cluster_id</i>	=, <, <=, >, >=	No	The master cluster ID.
	<i>master_cluster_name</i>	=	Yes	The master cluster name.
	<i>master_vdisk_id</i>	=, <, <=, >, >=	No	The master VDisk ID.
	<i>master_vdisk_name</i>	=	Yes	The master VDisk name.
	<i>aux_cluster_id</i>	=, <, <=, >, >=	No	The aux cluster ID.
	<i>aux_cluster_name</i>	=	Yes	The aux cluster name.
	<i>aux_vdisk_id</i>	=, <, <=, >, >=	No	The aux VDisk ID.
	<i>aux_vdisk_name</i>	=	Yes	The aux VDisk name.
	<i>primary</i>	=	No	The relationship primary. The following values are valid for <i>primary</i> : <ul style="list-style-type: none"> • master • aux
	<i>consistency_group_id</i>	=, <, <=, >, >=	No	The Metro Mirror consistency group ID.
	<i>consistency_group_name</i>	=	Yes	The Metro Mirror consistency group name.
	<i>state</i>	=	Yes	The relationship state. The following values are valid for <i>state</i> : <ul style="list-style-type: none"> • inconsistent_stopped • inconsistent_copying • consistent_stopped • consistent_synchronized • idling • idling_disconnected • inconsistent_disconnected • consistent_disconnected
<i>progress</i>	=, <, <=, >, >=	No	The progress of the initial background copy (synchronization) for the relationship.	

Table 27. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
rconsistgrp	<i>group_id or id</i>	=, <, <=, >, >=	No	The consistency group ID.
	<i>name</i>	=	Yes	The consistency group name.
	<i>master_cluster_id</i>	=, <, <=, >, >=	No	The master cluster ID.
	<i>master_cluster_name</i>	=	Yes	The master cluster name.
	<i>aux_cluster_id</i>	=, <, <=, >, >=	No	The aux cluster ID.
	<i>aux_cluster_name</i>	=	Yes	The aux cluster name.
	<i>primary</i>	=	No	The consistency group primary. The following values are valid for <i>primary</i> : <ul style="list-style-type: none"> • master • aux
	<i>state</i>	=	No	The consistency group state. The following values are valid for <i>state</i> : <ul style="list-style-type: none"> • inconsistent_stopped • inconsistent_copying • consistent_stopped • consistent_synchronized • idling • idling_disconnected • inconsistent_disconnected • consistent_disconnected • empty
<i>relationship_count</i>	=, <, <=, >, >=	No	The relationship count.	
user	<i>password</i>	=	No	Specifies if a password is associated with the user. The valid values are yes or no .
	<i>ssh_key</i>	=	No	Specifies if a Secure Shell (SSH) public key is associated with the user. The valid values are yes or no .
	<i>remote</i>	=	No	Specifies if the user authenticates to the cluster using a remote authentication service. The valid values are yes or no .
	<i>usergrp_id</i>	=, <, <=, >, >=	No	The ID of the user group.
	<i>usergrp_name</i>	=	Yes	The name of the user group.
usergrp	<i>role</i>	=	No	The role associated with all users that belong to this user group. The valid values are Monitor , CopyOperator , Service , Administrator , or SecurityAdmin .
	<i>remote</i>	=	No	Specifies if the user group is used to set the role of remote users. The valid values are yes or no .

Table 27. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
clusterip	<i>port_id</i>	=, <, <=, >, >=	No	The port ID. The valid values are 1 or 2 .
	<i>cluster_name</i>	=	Yes	The cluster name.
	<i>cluster_id</i>	=, <, <=, >, >=	No	The cluster ID.

Chapter 21. Overview of the dumps commands

| The **lsdump** command returns a list of dumps in a particular directory.

Dumps are contained in the following directory structure:

- /dumps
- /dumps/audit
- /dumps/cimom
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- | • /home/admin/upgrade
- | • /dumps/drive
- | • /dumps/enclosure

| Use the **lsdump** command with the optional **prefix** parameter to specify a directory. If you do not specify a directory, /dumps is used as the default. Use the optional *node_id_or_name* parameter to specify the node to list the available dumps for. If you do not specify a node, the available dumps on the configuration node are listed.

An audit log keeps track of action commands that are issued through an SSH session or from the management GUI. To list a specified number of the most recently audited commands, issue the **svctask catauditlog** command. To dump the contents of the audit log to a file on the current configuration node, issue the **svctask dumpauditlog** command. This command also clears the contents of the audit log.

Dumps contained in the /dumps/cimom directory are created by the CIMOM (Common Information Model Object Manager) that runs on the cluster. These files are produced during normal operations of the CIMOM.

Dumps that are contained in the /dumps/elogs directory are dumps of the contents of the error and event log at the time that the dump was taken. An error or event log dump is created by using the **svctask dumperrlog** command. This dumps the contents of the error or event log to the /dumps/elogs directory. If no file name prefix is supplied, the default errlog_ is used. The full default file name is errlog_NNNNNN_YYMMDD_HHMMSS, where NNNNNN is the node front panel name. If the command is used with the **-prefix** parameter, the prefix value is used instead of errlog.

Dumps contained in the /dumps/feature directory are dumps of the featurization log. A featurization log dump is created by using the **svctask dumpinternallog** command. This dumps the contents of the featurization log to the /dumps/feature directory to a file called feature.txt. Only one of these files exists, so every time the **svctask dumpinternallog** command is run, this file is overwritten.

| Dumps that are contained in the /dumps/iostats directory are dumps of the per-node I/O statistics for disks on the cluster. An I/O statistics dump is created by using the **svctask startstats** command. As part of this command, you can specify a time interval for the statistics to be written to the file; the default is 15 minutes. Every time the time interval is encountered, the I/O statistics that have been collected are written to a file in the /dumps/iostats directory. The file names that are used for storing I/O statistics dumps are Nm_stats_NNNNNN_YYMMDD_HHMMSS, Nv_stats_NNNNNN_YYMMDD_HHMMSS, Nn_stats_NNNNNN_YYMMDD_HHMMSS, and Nd_stats_NNNNNN_YYMMDD_HHMMSS, where NNNNNN is the node name for the MDisk, VDisk, node, or drive.

Dumps that are contained in the `/dumps/iotrace` directory are dumps of I/O trace data. The type of data that is traced depends on the options specified by the `svctask settrace` command. The collection of the I/O trace data is started by using the `svctask starttrace` command. The I/O trace data collection is stopped when the `svctask stoptrace` command is used. It is when the trace is stopped that the data is written to the file. The file name is `prefix_NNNNNN_YYMMDD_HHMMSS`, where *prefix* is the value entered for the `filename` parameter in the `svctask settrace` command, and `NNNNNN` is the node name.

Software upgrade packages are contained in the `/home/admin/upgrade` directory. These directories exist on every node in the cluster.

- | Dumps of support data from a disk drive are contained in the `/dumps/drive` directory. This data can help to identify problems with the drive, and does not contain any data that applications may have written to the drive.
- | Dumps from an enclosure or enclosures are contained in the `/dumps/enclosure` directory.

Dumps that are contained in the `/dumps` directory result from application abends. Such dumps are written to the `/dumps` directory. The default file names are `dump.NNNNNN.YYMMDD.HHMMSS`, where `NNNNNN` is the node front panel name. In addition to the dump file, there might be some trace files written to this directory that are named `NNNNNN.trc`.

Because files can only be copied from the current configuration node (using secure copy), you can issue the `svctask cpdumps` command to copy the files from a nonconfiguration node to the current configuration node.

Chapter 22. Information commands

The following commands enable you display specific types of SAN Volume Controller information.

Note: IDs are assigned at run-time by the system and cannot be relied upon to be the same after configuration restoration. Therefore, use object names instead of IDs whenever possible.

caterrlog (Deprecated)

The **caterrlog** command has been deprecated. Use the **lseventlog** command instead.

caterrlogbyseqnum (Deprecated)

The **caterrlogbyseqnum** command has been deprecated. Use the **lseventlog** command instead.

ls2145dumps (Deprecated)

- | **Attention:** The **svcinfol** **ls2145dumps** command is deprecated. Use the **svcinfol** **lsdumps** command to display a list of files in a particular dumps directory.

lscimomdumps (Deprecated)

- | **Attention:** The **svcinfol** **lscimomdumps** command is deprecated. Use the **svcinfol** **lsdumps** command to display a list of files in a particular dumps directory.

lscopystatus

Use the **lscopystatus** command to determine whether any file copies are currently in progress.

Syntax

```
➤— svcinfol — — lscopystatus — — [ -nohdr ] — — [ -delim delimiter ] — — ➤
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays an indicator that shows if a file copy is currently in progress. Only one file can be copied in the cluster at a time.

An invocation example

```
svcinfo lscopystatus
```

The resulting output

```
status
active
```

Iscluster

The `Iscluster` command returns a concise list or a detailed view of a cluster.

Syntax

```
➤➤ svcinfo -- lscluster -- [ -filtervalue -- attribute=value ] [ -nohdr ]
➤ [ -bytes ] [ -delim -- delimiter ] [ -filtervalue? ] [ cluster_id
  cluster_name ]
```

Parameters

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed. If a capacity is specified, the units must also be included.

Note: Some filters allow the asterisk character (*) when you enter the command. The following rules apply to the use of wildcard characters with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, you must enclose the filter entry within double quotation marks (""), as follows:

```
svcinfo lscluster -filtervalue "name=md*"
```

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-bytes

(Optional) Specifies that you want the report to display all capacities as bytes.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all

items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-filtervalue?

(Optional) displays a list of filters that can be applied against this view. The following filter attributes are valid for the **lscluster** command:

- **cluster_name**
- **cluster_unique_id**
- **id**
- **name**

cluster_id | cluster_name

(Optional) Specifies the name or ID of a cluster. When you use this parameter, the detailed view of the specific cluster is displayed and any value that you specified by the **-filtervalue** parameter is ignored. If you do not specify the *cluster_id | cluster_name* parameter, the concise view of all clusters that match the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

Description

This command displays a concise list or a detailed view of a cluster.

Table 28 provides the attribute values that can be displayed as output view data.

Table 28. Attribute values

Attribute	Possible Values
location	local, remote
statistics status	on, off
email_state	running, stopped, invalid
partnership	fully_configured, partially_configured_local, partially_configured_local_stopped, not present, fully_configured_stopped, fully_configured_remote_stopped, fully_configured_local_excluded, fully_configured_remote_excluded, fully_configured_exceeded
tier	Which tier information is being reported
tier_capacity	The total MDisk storage in the tier.
tier_free_capacity	The amount of MDisk storage in the tier that is unused.

Information about the remote cluster is reported by the **lscluster** command if the **mkpartnership** command has been issued from the local cluster to the remote cluster; for example, if the partnership has been at least partially established from the local cluster.

You can issue the **svcinfo lscluster** command to display a detailed view of the cluster.

```
svcinfo lscluster -delim : 10030a007e5
```

where **10030a007e5** is the cluster name.

Detailed view shows the fields described for remote clusters only; if the cluster **Location** is **local**, then **Partnership** and **Bandwidth** do not apply (and are not defined or provided). For a remote cluster, these fields indicate the following information:

- **Location:** remote

- **Partnership:**

- fully_configured**

- The **mkpartnership** command has been issued in both directions and the remote cluster is online and available.

- partially_configured_local**

- The **mkpartnership** command has only been issued from the local cluster to the remote cluster. The remote cluster is online and available for partnership.

- partially_configured_local_stopped**

- The **mkpartnership** command has only been issued from the local cluster to the remote cluster. The **chpartnership** command with the **stop** parameter has been issued from the local cluster, and the remote cluster is online and available. You need to issue the **chpartnership** command with the **start** parameter on the local cluster, and **mkpartnership** on the remote cluster.

- Not Present**

- The **mkpartnership** command has been issued from the local cluster to the remote cluster, and the remote cluster is not available. Either the remote cluster is offline, or it is not connected to the local cluster.

- fully_configured_stopped**

- The **mkpartnership** command has been issued in both directions and the remote cluster is online and available. The **chpartnership** command with the **stop** parameter has been issued from the local cluster.

- fully_configured__remote_stopped**

- The **mkpartnership** command has been issued in both directions and the remote cluster is online and available. The **chpartnership** command with the **stop** parameter has been issued from the remote cluster.

- fully_configured__local_excluded**

- The **mkpartnership** command has been issued in both directions. The local cluster has excluded the connection to the remote cluster due to too many problems, or either cluster in the partnership is unable to sustain the I/O workload for the Metro Mirror or Global Mirror relationships.

- fully_configured__remote_excluded**

- The **mkpartnership** command has been issued in both directions. The remote cluster has excluded the connection to the local cluster due to too many problems, or either cluster in the partnership is unable to sustain the I/O workload for the Metro Mirror or Global Mirror relationships.

- fully_configured__exceeded**

- There are too many clusters in the cluster network, and the partnership from the local cluster to the remote has been disabled. Refer to the 1710/1720 errors in the cluster error log at the local and remote clusters.

- **Bandwidth:** The bandwidth available on the intercluster link for background copy, in megabytes per second (MBps).

A concise invocation example

```
svcinfo lscluster -delim :
```

The concise resulting output

```
id:name:location:partnership:bandwidth:id_alias  
000002006420A162:cluster0:local:::000002006420A162
```

| An invocation example

```
| svcinfo lscluster cluster1
```

| **The resulting output**

```
| id 0000020061C15D90
| name cluster1
| location local
| partnership
| bandwidth
| total_mdisk_capacity 8.1TB
| space_in_mdisk_grps 8.1TB
| space_allocated_to_vdisks 8.00TB
| total_free_space 131.9GB
| statistics_status on
| statistics_frequency 15
| required_memory 0
| cluster_locale en_US
| time_zone 522 UTC
| code_level 6.1.0.0 (build 42.8.1006110000)
| FC_port_speed 2Gb
| console_IP 9.1.114.20:443
| id_alias 0000020061C15D90
| gm_link_tolerance 300
| gm_inter_cluster_delay_simulation 0
| gm_intra_cluster_delay_simulation 0
| email_reply
| email_contact
| email_contact_primary
| email_contact_alternate
| email_contact_location
| email_state stopped
| inventory_mail_interval 0
| total_vdiskcopy_capacity 8.00TB
| total_used_capacity 8.00TB
| total_overallocation 98
| total_vdisk_capacity 8.00TB
| cluster_ntp_IP_address
| cluster_isns_IP_address
| iscsi_auth_method none
| iscsi_chap_secret
| auth_service_configured no
| auth_service_enabled no
| auth_service_url
| auth_service_user_name
| auth_service_pwd_set no
| auth_service_cert_set no
| relationship_bandwidth_limit 25
| gm_max_host_delay 5
| tier_generic_ssd
| tier_capacity 0.00MB
| tier_free_capacity 0.00MB
| tier_generic_hdd
| tier_capacity 8.13TB
| tier_free_capacity 131.88GB
| email_contact2
| email_contact2_primary
| email_contact2_alternate
| total_allocated_extent_capacity 8.00TB
```

| **An invocation example for a cluster with two tiers**

```
| svcinfo lscluster china6
```

| **The resulting output**

```
| id 000002006220A5EE
| name china6
| location local
| partnership fully_configured_stopped
| bandwidth
```

```

| total_mdisk_capacity 474.1GB
| space_in_mdisk_grps 474.1GB
| space_allocated_to_vdisks 1.87GB
| total_free_space 472.2GB
| statistics_status on
| statistics_frequency 15
| required_memory 0
| cluster_locale en_US
| time_zone 325 Etc/GMT+1
| code_level 6.1.0.0 (build 46.4.1008240000)
| FC_port_speed 2Gb
| console_IP 9.71.43.210:443
| id_alias 000002006220A5EE
| gm_link_tolerance 300
| gm_inter_cluster_delay_simulation 0
| gm_intra_cluster_delay_simulation 0
| email_reply:fred@mycompany.com
| email_contact:Fred Higgins
| email_contact_primary:01202 123456
| email_contact_alternate:44-202-876543-4455
| email_contact_location:London Thames Bank
| email_state:running
| inventory_mail_interval:0
| total_vdiskcopy_capacity 30.58GB
| total_used_capacity 1.28GB
| total_overallocation 6
| total_vdisk_capacity 20.27GB
| cluster_ntp_IP_address
| cluster_isns_IP_address
| iscsi_auth_method none
| iscsi_chap_secret
| auth_service_configured:yes
| auth_service_enabled:yes
| auth_service_url:https://1.2.3.4/login
| auth_service_user_name:secadmin
| auth_service_pwd_set:yes
| auth_service_cert_set:yes
| relationship_bandwidth_limit 25
| gm_max_host_delay 5
| tier_generic_ssd
| tier_capacity 270.91GB
| tier_free_capacity 270.03GB
| tier_generic_hdd
| tier_capacity 203.19GB
| tier_free_capacity 201.94GB
| email_contact2:barry
| email_contact2_primary:6347
| email_contact2_alternate:6589
| total_allocated_extent_capacity 2.97GB

```

Isclustercandidate

The **Isclustercandidate** command lists the clusters that are available for setting up a partnership with the local cluster. This is a prerequisite for creating intercluster Metro or Global Mirror relationships.

Syntax

```

▶▶ svcinfo — — Isclustercandidate — — [ -nohdr ] — — [ -delim — delimiter ] — —▶▶

```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays a list of clusters that are available as candidate partner clusters to form a Metro Mirror or Global Mirror partnership between two clusters.

Output from the command shows the cluster ID, name, and configured status of the remote candidate cluster. The remote candidate cluster forms a partnership with the local cluster when you use the **svctask mkpartnership** command. The remote cluster shows the partnership status as `partially_configured_local_stopped` or `partially_configured_local` when you use the **svcinfolcluster** command. The **svcinfolclustercandidate** command displays the configured status of those remote clusters that have formed a partnership with the local cluster.

An invocation example

```
svcinfolclustercandidate
```

The resulting output

id	configured	cluster_name
0000010034E0F430	no	ldcluster26

lsclusterip

The **lsclusterip** command returns a list of the cluster management IP addresses configured for each port.

Syntax

```
svcinfolclusterip -nohdr -delim delimiter -filtervalue attribute=value -filtervalue? cluster_id cluster_name
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed. If a capacity is specified, the units must also be included.

Note: Some filters allow the asterisk character (`*`) when you enter the command. The following rules apply to the use of wildcard characters with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (`*`).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, you must enclose the filter entry within double quotation marks (`"`), as follows:

```
svcinfolclusterip -filtervalue "cluster_name=md*
```

-filtervalue?

(Optional) displays a list of filters that can be applied against this view. The following filter attributes are valid for the **lsclusterip** command:

- **port_id**
- **cluster_name**
- **cluster_id**

cluster_id | cluster_name

(Required) Specifies the name or ID of a cluster.

Description

This command displays a list of the cluster management IP addresses configured for each port.

A concise invocation example

```
svcinfolclusterip -delim ,
```

The concise resulting output

```
cluster_id,cluster_name,location,port_id,IP_address,subnet_mask,
gateway,IP_address_6,gateway_6,prefix_6
000002006CC0B71A,c11,local,1,192.168.1.2,DHCP,255.255.255.0,192.168.1.1,
2001:0db8:85a3:0000:0000:8a2e:0370:7334,2001:0db8:85a3:0000:0000:8a2e:0370:7334,
2001:0db8:85a3:0000:0000:8a2e:0370:7334,64
000002006CC0B71A,c11,local,2,192.168.1.2,DHCP,255.255.255.0,192.168.1.1,
2001:0db8:85a3:0000:0000:8a2e:0370:7334,2001:0db8:85a3:0000:0000:8a2e:0370:7334,
2001:0db8:85a3:0000:0000:8a2e:0370:7334,64
000002006CC0B7110,c12,remote,1,192.168.1.2,DHCP,255.255.255.0,192.168.1.1,
2001:0db8:85a3:0000:0000:8a2e:0370:7334,2001:0db8:85a3:0000:0000:8a2e:0370:7334,
2001:0db8:85a3:0000:0000:8a2e:0370:7334,64
000002006CC0B7110,c12,remote,2,192.168.1.2,DHCP,255.255.255.0,192.168.1.1,
2001:0db8:85a3:0000:0000:8a2e:0370:7334,2001:0db8:85a3:0000:0000:8a2e:0370:7334,
2001:0db8:85a3:0000:0000:8a2e:0370:7334,64
```

A detailed invocation example

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-filtervalue?

(Optional) Displays the valid filter attributes. The following filter attributes for the **svcinfo lscontroller** command are valid:

- `controller_id`
- `id`

controller_id | *controller_name*

(Optional) Specifies the name or ID of a controller. When you use this parameter, the detailed view of the specific controller is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *controller_id* | *controller_name* parameter, the concise view displays all controllers matching the filtering requirements that are specified by the **-filtervalue** parameter.

Description

This command returns a concise list, or a detailed view, of controllers visible to the cluster.

The following values are applicable to the data in the output views:

degraded no, yes

To differentiate the name of a storage controller from the name shown on the cluster, list the storage controllers by issuing the **svcinfo lscontroller** command. Record the controller name or ID for the controller that you want to determine. For the controller in question, issue the **svcinfo lscontroller *controller name* | *id*** command, where *controller name* | *id* is the controller name or ID. Record the worldwide node name (WWNN) for the controller. You can use the WWNN to determine the actual storage controller by launching the native controller user interface, or by using the command line tools it provides to verify the actual controller that has the WWNN.

Notes:

1. The *mdisk_link_count* value is the number of MDisks currently using this controller as their active controller.
2. The *max_mdisk_link_count* value is the highest value that the *mdisk_link_count* has reached since it was last reset to the *mdisk_link_count* value. This value is reset by specific maintenance procedures and when the cluster error log is cleared.
3. The controller *path_count* value is the sum of the MDisk *path_count* values for MDisks currently using this controller port as their active controller port.
4. The controller *max_path_count* value is the highest value that the controller *path_count* has reached since it was last reset to the *path_count* value. This value is reset by specific maintenance procedures and when the cluster error log is cleared.
5. The *allow_quorum* value shows whether the controller is currently enabled to support quorum disks. Quorum support is either enabled or disabled depending on the controller hardware type.
6. The *ctrl_s/n* value is the controller serial number. This data comes from vendor-controlled sources and might not be available.

A concise invocation example

```
svcinfo lscontroller -delim :
```

The concise resulting output

```
id:controller_name:ctrl_s/n:vendor_id:product_id_low:product_id_high
7:controller7:3EK0J5Y8:SEAGATE :ST373405:FC
8:controller8:3EK0J6CR:SEAGATE :ST373405:FC
9:controller9:3EK0J4YN:SEAGATE :ST373405:FC
10:controller10:3EK0GKGH:SEAGATE :ST373405:FC
11:controller11:3EK0J85C:SEAGATE :ST373405:FC
12:controller12:3EK0JBR2:SEAGATE :ST373405:FC
13:controller13:3EKYNJF8:SEAGATE :ST373405:FC
14:controller14:3EK0HVTM:SEAGATE :ST373405:FC
```

A detailed invocation example

```
svcinfolcontroller -delim = 7
```

The detailed resulting output

```
id=7
controller_name=controller7
WWNN=20000004CF2412AC
mdisk_link_count=1
max_mdisk_link_count=1
degraded=no
vendor_id=SEAGATE
product_id_low=ST373405
product_id_high=FC
product_revision=0003
ctrl_s/n=3EK0J5Y8
allow_quorum=no
WWPN=22000004CF2412AC
path_count=1
max_path_count=1
WWPN=21000004CF2412AC
path_count=0
max_path_count=0
```

Iscontrollerdependentvdisks

The `Iscontrollerdependentvdisks` command lists the VDisks that are dependent on the specified controller.

Syntax

```
▶▶ svcinfol — — Iscontrollerdependentvdisks — — [ controller_id_list | controller_name_list ] ▶▶
```

Parameters

controller_id_list | *controller_name_list*

Specifies one or more controller IDs, controller names, or both. Separate multiple controllers using the colon character (:).

Description

The `Iscontrollerdependentvdisks` command lists the VDisks that are dependent on the status of the specified controllers. If a controller goes offline, the dependent VDisks also go offline. Before taking a controller offline for maintenance, you can use the command to ensure that you do not lose access to any VDisks.

If you have multiple controllers configured as a single subsystem, you must specify all of the controllers in the subsystem, using a single command invocation.

The **lscontrollerdependentvdisks** command also checks for quorum disks on the specified controller list. If any quorum disks are on the specified controller list, the command returns an error. All quorum disks must be moved before performing any maintenance. After moving quorum disks, reissue the command to list the dependent VDisks.

Note: The command lists the VDisks that are dependent on the controllers at the time the command is run; subsequent changes to your system require rerunning the command.

An invocation example

```
svcinfo lscontrollerdependentvdisks controller0
```

The resulting output

```
vdisk_id vdisk_name
0 vdisk0
1 vdisk1
2 vdisk2
```

lscurrentuser

Use the **lscurrentuser** command to display the name and role of the logged-in user.

Syntax

```
▶▶ svcinfo — — lscurrentuser — — [ -nohdr ] [ -delim delimiter ] ▶▶
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays the name and role of the current user.

An invocation example

```
svcinfo lscurrentuser
```

The resulting output

```
name superuser
role SecurityAdmin
```

lsdiscoverystatus

Use the **lsdiscoverystatus** command to determine whether a discovery operation is in progress.

Syntax

```
➤— svcinfo — — lsdiscoverystatus — — [ -nohdr ] — — [ -delim — delimiter ] — — ➤
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

| This command displays the state of all discoveries in the cluster. During discovery, the system updates
| the drive and MDisk records. You must wait until the discovery has finished and is inactive before you
| attempt to use the system. This command displays one of the following results:

active There is a discovery operation in progress at the time that the command is issued.

inactive

There are no discovery operations in progress at the time that the command is issued.

| If the Fibre Channel functions are used only to enable the nodes to cluster, then the Fibre Channel line
| will not be displayed in the **svcinfo lsdiscoverystatus** command. The *fc_fabric* line will only appear if
| there is at least one Fibre Channel controller.

An invocation example

```
svcinfo lsdiscoverystatus -delim :
```


The resulting output

```
id:scope:IO_group_id:IO_group_name:status  
0:fc_fabric:::active  
1:sas_iogrp:0:io_grp0:inactive  
3:sas_iogrp:2:io_grp2:active
```

lsdumps

| Use the **lsdumps** command to display a list of files in a particular dumps directory on one of the nodes
| in the cluster.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, then the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-prefix *directory_name*

(Optional) Specifies the name of the directory to list files for. The default is the `/dumps` directory. Valid directory names:

- `/dumps`
- `/dumps/audit`
- `/dumps/cimom`
- `/dumps/elogs`
- `/dumps/feature`
- `/dumps/iostats`
- `/dumps/iotrace`
- `/dumps/mdisk`
- `/home/admin/upgrade`
- `/dumps/drive`
- `/dumps/enclosure`

node_id_or_name

(Optional) Specifies the node ID or name to list the available dumps for. If you do not specify a node, the available dumps on the configuration node are listed.

Description

This command displays a list of files detected by a node. You can specify the name of the directory to list files for, and the node ID or name. If you do not specify a directory, the `/dumps` directory is used.

The files are listed in order of time created, with the oldest files listed first.

An invocation example

| To list the files in /dumps on the configuration node:

```
| svcinfo lsdumps
```

| The resulting output

```
| id  filename
| 0   svc.config.cron.bak_node1
| 1   svc.config.backup.xml_node1
| 2   recover.110584.100116.035201
| 3   dump.110584.100118.051550
| 4   ethernet.aaabbbX-1.trc
```

| An invocation example

| To list the error log files on node 2:

```
| svcinfo lsdumps -prefix /dumps/eologs 2
```

| The resulting output

```
| id  filename
| 0   errlog_110584_090624_200258
| 1   errlog_110584_090717_231023
```

lsemailer

The **lsemailer** command returns a concise list or a detailed view of e-mail servers that are configured on the cluster.

Syntax

```
▶▶ svcinfo — — lsemailer — — [ -nohdr ] — — [ -delim — delimiter ] — —————▶
▶ [ email_server_name ] —————▶
  [ email_server_id ] —————▶
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

email_server_name | *email_server_id*

(Optional) Specifies the name or ID of an existing e-mail server that must be listed.

Description

Use this command to display a concise list or a detailed view of e-mail servers that are configured on the cluster.

A concise invocation example

```
svcinfolsemailserver -delim :
```

The concise resulting output

```
id:name:IP_address:port
0:emailserver0:192.135.60.3:25
1:emailserver1:192.135.60.4:25
2:emailserver2:192.135.60.5:25
```

A detailed invocation example

```
svcinfolsemailserver email0
```

The detailed resulting output

```
id 0
name emailserver0
IP_address 192.135.60.3
port 25
```

lsemailuser

The **lsemailuser** command generates a report that lists the e-mail event notification settings for all e-mail recipients, an individual e-mail recipient, or a specified type (local or support) of e-mail recipient.

Syntax

```
➤➤ svcinfolsemailuser -type support | local -delim delimiter id_or_name
```

Parameters

-type support | local

(Optional) Specifies the types of e-mail recipients you want to view, either customer or support based as determined by the following definitions:

support

Address of the support organization that provides vendor support.

local All other addresses.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, a colon separates all items of data in a concise view; the spacing of columns does not occur. In a detailed view, the data is separated from its header by a colon.

id_or_name

(Optional) Specifies the user ID or user name of the e-mail event recipient for whom you want to see the e-mail notification settings.

Description

When you issue this command, a report is displayed that lists the e-mail event notification settings for all e-mail recipients, an individual e-mail recipient, or a specified type (local or support) of e-mail recipient. The concise and detailed views report the same information.

An invocation example

The following command lists information for all e-mail recipients using the e-mail event notification facility, in a concise view:

```
svcinfolsemailuser -delim :
```

The resulting output

```
id:name:address:user_type:error:warning:info:inventory
1:Support:callhome1@de.ibm.com:support:on:off:off:off
2:Fred:fred_house@my_company.co.uk:local:on:on:on:off
3:Log:our_log@my_company.co.uk:local:on:on:on:on
```

Iserrlogbyfcconsistgrp (Deprecated)

The `Iserrlogbyfcconsistgrp` command has been deprecated. Use the `lseventlog` command instead.

Iserrlogbyfcmap (Deprecated)

The `Iserrlogbyfcmap` command has been deprecated. Use the `lseventlog` command instead.

Iserrlogbyhost (Deprecated)

The `Iserrlogbyhost` command has been deprecated. Use the `lseventlog` command instead.

Iserrlogbyiogrp (Deprecated)

The `Iserrlogbyiogrp` command has been deprecated. Use the `lseventlog` command instead.

Iserrlogbymdisk (Deprecated)

The `Iserrlogbymdisk` command has been deprecated. Use the `lseventlog` command instead.

Iserrlogbymdiskgrp (Deprecated)

The `Iserrlogbymdiskgrp` command has been deprecated. Use the `lseventlog` command instead.

Iserrlogbynode (Deprecated)

The `Iserrlogbynode` command has been deprecated. Use the `lseventlog` command instead.

Iserrlogbyrcconsistgrp (Deprecated)

The `Iserrlogbyrcconsistgrp` command has been deprecated. Use the `lseventlog` command instead.

Iserrlogbyrelationship (Deprecated)

The `Iserrlogbyrelationship` command has been deprecated. Use the `Iseventlog` command instead.

Iserrlogbyvdisk (Deprecated)

The `Iserrlogbyvdisk` command has been deprecated. Use the `Iseventlog` command instead.

Iserrlogdumps (Deprecated)

| **Attention:** The `svcinfolerrlogdumps` command is deprecated. Use the `svcinfolsdumps` command to display a list of files in a particular dumps directory.

Iseventlog

| Use the `Iseventlog` command to display a concise view of the system event log, or a detailed view of one entry from the log.

Syntax

```
| ▶▶ svcinfo -- lseventlog [-alert yes|no] [-message yes|no] [-config yes|no]
|
| [-monitoring yes|no] [-expired yes|no] [-fixed yes|no] [-count entry_limit]
|
| [-order date|severity] [sequence_number]
|
```

Parameters

| **-alert**
| Includes events with *alert* status.

| **-message**
| Includes events with *message* status.

| **-config**
| Includes events with *config* status.

| **-monitoring**
| Includes events with *monitoring* status.

| **-expired**
| Includes events with *expired* status.

| **-fixed**
| Includes fixed events.

| **-count**
| Indicates the maximum number of events to display.

| **-order** *date|severity*
| Indicates what order the events should be in. Ordering by *date* displays the oldest events first.
| Ordering by *severity* displays the events with the highest severity first. If multiple events have the
| same severity, then they are ordered by date, with the oldest event being displayed first.

| The following list shows the order of severity, starting with the most severe:

- | 1. Unfixed alerts (sorted by error code; the lowest error code has the highest severity)

2. Unfixed messages
3. Monitoring events (sorted by error code; the lowest error code has the highest severity)
4. Expired events
5. Configuration events
6. Fixed alerts and messages

sequence_number

Indicates if the command should display a full view of the event.

Description

This command displays a concise view of the system event log, or a detailed view of one entry from the log. You can sort the events and entries by severity or age.

Table 29 provides the attribute values that can be displayed as output view data.

Table 29. *lseventlog* output

Attribute	Description	Value
<i>sequence_number</i>	Sequence number of the event.	Numeric 0-8000000
<i>first_timestamp</i>	When the event was added to the log.	YYMMDDHHMMSS
<i>first_timestamp_epoch</i>	When the event was added to the log (in seconds) since the epoch	Numeric 32-bit
<i>last_timestamp</i>	When the event was most recently updated.	YYMMDDHHMMSS
<i>last_timestamp_epoch</i>	When the event was most recently updated (in seconds) since the epoch.	Numeric 32-bit
<i>object_type</i>	The type of the object the event is logged against.	<ul style="list-style-type: none"> • <i>mdisk</i> • <i>mdiskgrp</i> • <i>vdisk</i> (or <i>vdisk copy</i>) • <i>node</i> • <i>host</i> • <i>io_grp</i> (<i>iogroup</i> in dumperrlog) • <i>fc_consist_grp</i> (<i>fcgrp</i> in dumperrlog) • <i>rc_consist_grp</i> (<i>rcgrp</i> in dumperrlog) • <i>fc_map</i> (<i>fcmap</i> in dumperrlog; <i>flash</i> in catterrlog) • <i>rc_relationship</i> (<i>rcmap</i> in dumperrlog; <i>remote</i> in catterrlog) • <i>cluster</i> • <i>controller</i> (<i>device</i> in catterrlog and dumperrlog) • <i>quorum</i> • <i>migrate</i> • <i>email_server</i> (<i>email server</i> in catterrlog and dumperrlog) • <i>enclosure</i> • <i>drive</i>

Table 29. lseventlog output (continued)

Attribute	Description	Value
object_id	The ID of the object the event is logged against.	Numeric 64-bit. Displayed in decimal for all object types except clusters. Clusters are in hexadecimal. Blank for events with object_type <i>cluster</i> .
object_name	The name of the object the event is logged against.	Object name format; blank if the object was deleted or does not have a name.
copy_id	The VDisk copy ID the event is logged against.	0-1; blank if not a vdiskcopy event.
reporting_node_id	The ID of the node that reported the event.	Numeric 64-bit; blank if the event was reported by the cluster.
reporting_node_name	The name of the node that reported the event.	Object name format; blank if the node was deleted or if the event was reported by the cluster.
root_sequence_number	Sequence number of the root, causal event.	Numeric, 1-8000000; blank if there is no root.
event_count	The number of reported events that have been combined into this event.	Numeric 32-bit.
status	The event category.	<ul style="list-style-type: none"> • alert • message • config • monitoring • expired
fixed	Indicates whether the event was marked fixed (for an alert) or read (for a message).	<ul style="list-style-type: none"> • yes • no (for events that cannot be fixed, or are not fixed)
auto_fixed	Indicates whether event was marked fixed by the code.	<ul style="list-style-type: none"> • yes • no (for events that cannot be fixed, or are not fixed)
notification_type	The type of event notification.	<ul style="list-style-type: none"> • error • warning • informational • none
event_id	The event ID.	6-digit numeric
event_id_text	The description associated with the event ID.	Text, max 200 bytes. Appears in the CLI requested language.
error_code	The error code associated with this event.	4-digit numeric; blank if there is no error code.
error_code_text	The description associated with the error code.	Text (maximum of 200 bytes); blank if there is no error code. Appears in CLI requested language.
description	The description associated with the event. If the event has an error code, this value is the same as the error_code_text field; otherwise, it is the same as the event_id_text field.	Text (maximum of 200 bytes).

Table 29. lseventlog output (continued)

Attribute	Description	Value
sense1	Sixteen bytes of hex-encoded sense data. The least significant byte is on the left.	Sixteen two-character hex numbers separated by spaces.
sense2		
sense3		
sense4		
sense5		
sense6		
sense7		
sense8		

An invocation example

This example shows events in January 2010:

```
lseventlog -filtervalue last_timestamp>=100101000000:last_timestamp<100201000000
```

This example shows all unfixed 1065 errors, in order of occurrence:

```
lseventlog -filtervalue error_code=1065:fixed=no
```

This example lists the most critical event:

```
lseventlog -order severity -count 1
```

This example shows the concise view:

```
lseventlog
sequence_number:last_timestamp:object_type:object_id:object_name:copy_id:
status:fixed:event_id:error_code:description
400:100106132413:vdisk:2:my_vdisk:1:alert:no:060001:1865:
Space Efficient Virtual Disk Copy offline due to insufficient space
401:100106140000:cluster::ldcluster-2::message:no:981001:
:Cluster Fabric View updated by fabric discovery
```

This example shows the full view:

```
lseventlog 400
sequence_number 400
first_timestamp 100106132413
first_timestamp_epoch 1256025652
last_timestamp 100106132413
last_timestamp_epoch 1256025652
object_type vdisk
object_id 2
object_name my_vdisk
copy_id 1
reporting_node_id 2
reporting_node_name node2
root_sequence_number
event_count 1
status alert
fixed no
auto_fixed no
notification_type warning
event_id 060001
event_id_text Space Efficient Virtual Disk Copy offline due to insufficient space
error_code 1865
error_code_text Space Efficient Virtual Disk Copy offline due to insufficient space
sense1 00 00 00 00 00 00 00 00 00 00 00 00 01 23 42 AB
```

```

| sense2 00 00 00 01 20 A3 14 22 01 00 00 00 00 02 00
| sense3 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
| sense4 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
| sense5 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
| sense6 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
| sense7 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
| sense8 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

```

lsfabric

The **lsfabric** command generates a report that displays the fibre-channel connectivity between nodes, controllers and hosts.

Syntax

```

➔— svcinfo — — lsfabric —————➔
|
| -node — node_id_or_name — [ -port — port_id ]
|
| -wwpn — wwpn
|
| -host — host_id_or_name
|
| -controller — controller_id_or_name
|
| -cluster — cluster_id_or_name

```

Parameters

-node *node_id_or_name*

(Optional) Displays the output for all ports for the specified node. The only parameter that you can specify with the **-node** parameter is the **-port** parameter.

-port *port_id*

(Optional) Displays a concise view of all WWPNs that are logged into the specified port ID and node. The **-port** parameter must be specified with only the **-node** parameter. A valid *port_id* value is a number from 1 - 4 that specifies the port number in the vital product data (VPD) or the hexadecimal WWPN of the local port.

-wwpn *wwpn*

(Optional) Displays a list of all ports that have a login to the specified WWPN. You cannot use the **-wwpn** parameter with any other parameter.

-host *host_id_or_name*

(Optional) Specifies a host name or ID. Issuing the **lsfabric** command with the **-host** parameter is equivalent to issuing the **svcinfo lsfabric -wwpn *wwpn*** command for every configured WWPN of the specified host. For example, a host with two ports that are zoned to one port of every node in a eight-node cluster produces 16 lines of output. You cannot use the **-host** parameter with any other parameter.

-controller *controller_id_or_name*

(Optional) Specifies a controller ID or name. You cannot use the **-controller** parameter with any other parameter in this command. Issuing the **lsfabric** command with the **-controller** parameter is equivalent to issuing the **svcinfo lsfabric -wwpn *wwpn*** command for every configured WWPN of the specified controller. For example, a controller with 4 ports connected to a 8 node cluster with 2 counterpart SANs produces 64 lines of output.

-cluster *cluster_id_or_name*

(Optional) Specifies a cluster ID or name. You cannot use the **-cluster** parameter with any other parameter. Issuing the **lsfabric** command with the **-cluster** parameter is equivalent to issuing the **svcinfo lsfabric -wwpn *wwpn*** command for every known WWPN in the specified cluster. Output is sorted by remote WWPNs and then cluster WWPNs. This parameter can be used to check the state of connections within the local cluster or between the local and remote cluster. When the local cluster ID

or name is specified, each node-to-node connection is listed twice: once from each end. For example, an eight-node cluster with two counterpart SANs produces eight nodes, multiplied by seven other nodes, multiplied by two SANs, multiplied by four point-to-point logins, equals 448 lines of output.

| **Note:** The cluster must be configured in a remote copy partnership with the local cluster; it must
| appear in the **svcinfolcluster** view.

Description

The **lsfabric** command can be issued with any of the parameters to display a limited subset of information. If the command is issued without any parameters, it provides output for every node.

Values for the **Type** and **State** columns are:

state active

The meaning of this value depends on the object that it applies to, as follows:

- **host or controller:** SCSI commands were issued within the last 5 minutes.
- **node:** node ports can see other ports.

state inactive

No transactions have completed within the last 5 minutes.

Note: It can take up to 10 seconds after a command for a controller port to change from inactive to active. It can take up to 5 minutes after a command for a host port to change from inactive to active.

type One of the following values is displayed:

- **host**
- **node**
- **controller**
- **unknown**

You can issue this command to view all the information about the connections that are available to your cluster.

An invocation example

```
svcinfolfabric -delim :
```

The resulting output Each row of output contains the following colon-separated columns:

```
remote_wwpn:remote_nportid:id:node_name:local_wwpn:  
local_port:local_nportid:state:name:cluster_name:type
```


lsfcconsistgrp

The **lsfcconsistgrp** command returns a concise list or a detailed view of FlashCopy consistency groups that are visible to the cluster. This information is useful for tracking FlashCopy consistency groups.

The list report style can be used to obtain two styles of report:

- A list containing concise information about all of the FlashCopy consistency groups on a cluster. (Each entry in the list corresponds to a single FlashCopy consistency group.)
- The detailed information about a single FlashCopy consistency group.

Syntax



Parameters

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk character (*).
- The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- When you use a wildcard, surround the filter entry with double quotation marks (""), as follows:

```
svcinfo lsfconsistgrp -filtervalue "name=md*"
```

-nohdr

(Optional) By default, headings are displayed for each item of data in a concise view. The **-nohdr** parameter suppresses the display of these headings. Detailed view is not valid for this command.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, the headers are displayed, and the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | object_name

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter returns an error message. If you do not specify the *object_id* or *object_name* parameter, the concise view of all objects matching the filtering requirements that is specified by the **-filtervalue** parameter are displayed.

-filtervalue?

(Optional) Displays the list of valid filter attributes in the report. The valid filter attributes for the **svcinfo lsfconsistgrp** command are:

- **name**
- **id**
- **status**
- **FC_group_id**

Description

This command returns a concise list or a detailed view of FlashCopy consistency groups that are visible to the cluster.

The following list provides values of the *status* attribute that are displayed as data in the output views:

status **empty, idle_or_copied, preparing, prepared, copying, stopped, suspended, stopping**

A concise invocation example

```
svcinfolsfcconsistgrp -delim :
```

The concise resulting output

```
id:name:status
1:ffccg0:empty
2:ffccg1:idle_or_copied
3:ffccg2:idle_or_copied
```

A detailed invocation example

```
svcinfolsfcconsistgrp -delim : 1
```

The detailed resulting output

```
id:1
name:ffccg0
status:empty
```

A detailed invocation example

```
svcinfolsfcconsistgrp -delim : fccstgrp0
```

The detailed resulting output

```
id:1
name:FCcgrp0
status:idle_or_copied
autodelete:off
FC_mapping_id:0
FC_mapping_name:fcmap0
FC_mapping_id:1
FC_mapping_name:fcmap1
```

lsfcmap

The **lsfcmap** command generates a list containing concise information about all of the FlashCopy mappings that are visible to the cluster, or detailed information for a single FlashCopy mapping.

Syntax

```
▶▶ svcinfo — lsfcmap — [ -filtervalue — attribute=value ] [ -nohdr ]
▶ [ -delim — delimiter ] [ -filtervalue? ] [ object_id ] [ object_name ]
```

Parameters

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is the asterisk (*), which must be used as the first or last character in the string.
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:

```
svcinfolsfcm -filtervalue "name=md*"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** *attribute=value* parameter:

- name
- id
- source_vdisk_id
- source_vdisk_name
- target_vdisk_id
- target_vdisk_name
- group_name
- group_id
- status
- copy_rate
- FC_mapping_name
- FC_id
- partner_FC_id
- partner_FC_name
- restoring

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **delim** parameter overrides this behavior. Valid input for the **delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | object_name

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **filtervalue** parameter is ignored. If you do not specify the **object_ID** or **object_name** parameter, the concise view of all objects matching the filtering requirements that is specified by the **filtervalue** parameter are displayed.

Description

This command returns a concise list or a detailed view of FlashCopy mappings that are visible to the cluster.

The following list shows attribute values that can be displayed as output view data:

status idle_or_copied, preparing, prepared, copying, stopped, suspended or stopping

start_time

Displays the time that the copy was last started. It is in the format *YYMMDDHHMMSS*. If a copy has not been started, a blank line is displayed.

A concise invocation example

```
svcinfc lsfcmap -delim :
```

The concise resulting output

```
id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:target_vdisk_name:
group_id:group_name:status:progress:copy_rate:clean_progress:incremental:
partner_FC_id:partner_FC_name:restoring
0:fcmap0:0:vdisk0:1:vdisk1:0:fccstgrp0:idle_or_copied:0:50:0:on:2:fcmap2:no
1:fcmap1:2:vdisk2:3:vdisk3:0:fccstgrp0:idle_or_copied:0:0:100:off:::no
2:fcmap2:1:vdisk1:0:vdisk0:0:fccstgrp1:idle_or_copied:0:0:100:off:0:fcmap0:no
```

A detailed invocation example

```
svcinfc lsfcmap -delim : 0
```

The detailed resulting output

```
id:0
name:fcmap0
source_vdisk_id:63
source_vdisk_name:vdisk63
target_vdisk_id:57
target_vdisk_name:vdisk57
group_id:
group_name:
status:idle_or_copied
progress:0
copy_rate:0
start_time:
dependent_mappings:0
autodelete:off
clean_progress:100
clean_rate:50
incremental:off
difference:100
grain_size:256
IO_group_id:1
IO_group_name:io_grp1
partner_FC_id:
partner_FC_name:
restoring:no
```

lsfcmapcandidate

The **lsfcmapcandidate** command lists all of the VDisks that are associated with fewer than 256 FlashCopy mappings.

Syntax

```
►► svcinfo — — lsfcmapcandidate — [ -nohdr ] [ -delim — delimiter ]
```

Parameters

-nohdr

(Optional) By default, the heading is displayed for the column of data in a concise style view, and for the item of data in a detailed style view. The **-nohdr** parameter suppresses the display of the heading.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, a colon character (:) separates all items of data in a concise view; the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command returns a list of VDisks that are associated with fewer than 256 FlashCopy mappings.

An invocation example

```
svcinfo lsfcmapcandidate
```

The resulting output

```
id  
2  
3  
4
```

lsfcmapprogress

The **lsfcmapprogress** command returns the progress of the background copy of a FlashCopy mapping. This is displayed as a percentage completed value.

Syntax

```
►► svcinfo — — lsfcmapprogress — [ -nohdr ] [ -delim — delimiter ]
```

```
► [ fcmapping_id ]  
  [ fcmapping_name ]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one byte character. If you enter `-delim :` on the command line, the data is separated from its header by a colon character (:).

fcmap_id | *fcmap_name*

(Required) Specifies that you want the report to display the progress of the background copy for the designated FlashCopy mapping.

Description

This command reports a percentage for the progress of the background copy being done on the specified FlashCopy mapping.

An invocation example

```
svcinfolsfcmapprogress 0
```

The resulting output

```
id          progress
0           0
```

lsfcmdependentmaps

The `lsfcmdependentmaps` command displays all the FlashCopy mappings that are dependent on the user specified mapping.

Syntax

```
▶▶ svcinfolsfcmdependentmaps [-nohdr] [-delim delimiter]
▶ [ fc_id | fc_name ]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

fc_id | *fc_name*

(Required) Specifies the name or ID of the FlashCopy mapping to list the dependent maps for.

Description

This command returns a list of dependent FlashCopy mappings. This command can be used to determine the list of FlashCopy mappings that would also stop if you stopped a mapping using the **-force** parameter.

There is a `dependent_mapping_count` field in the FlashCopy map detailed view (displayed when you process the **lsfcmap** command) that you can use as an indicator as to whether there are any dependent mappings in progress. If the count is zero, there are no dependent copies.

Note: If a period of time elapses between the time you process the **lsfcmap** command and the **lsfcmapdependentmaps** command, there could be a difference between the actual number of dependent mappings being processed and the number that was reported by the **lsfcmap** command.

A concise invocation example

```
svcinfolsfmapdependentmaps -delim : 2
```

The resulting output

```
fc_id:fc_name
1:fcmap1
3:fcmap3
```

lsfeaturedumps (Deprecated)

- | **Attention:** The **svcinfolsffeaturedumps** command is deprecated. Use the **svcinfolsdumps** command to display a list of files in a particular dumps directory.

lsfreeextents

The **lsfreeextents** command lists the number of free extents that are available on a specified MDisk.

Syntax

```
➤— svcinfolsfreeextents — [ -nohdr ] [ -delim delimiter ] —————➤
|
| [ mdisk_id ] —————➤
| [ mdisk_name ]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

mdisk_id | *mdisk_name*

(Required) Specifies the ID or the name of the MDisk for which you want to know the number of free extents.

Description

This command displays a count of the number of free extents on the specified MDisk.

An invocation example

```
svcinflsfreeextents 2
```

The resulting output

```
id 2
number_of_extents 4372
```

lshbaportcandidate

The **lshbaportcandidate** command lists all of the unconfigured, logged-in host bus adapter (HBA) ports. This information is used to find open HBA ports.

Syntax

```
svcinflshbaportcandidate [-nohdr] [-delim delimiter]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command returns a list of unconfigured, logged in HBA ports.

Note: The **svcinflshbaportcandidate** command presents a list of host HBA ports that are logged in to nodes. However, there are situations when the information that is presented might include host HBA ports that are no longer logged in or even part of the SAN fabric. For example, a host HBA port is unplugged from a switch but **svcinflshbaportcandidate** still shows the WWPN that is logged in to all nodes. If this occurs, the incorrect entry is removed when another device is plugged in to the same switch port that previously contained the removed host HBA port.

An invocation example

```
svcinflshbaportcandidate
```

The resulting output

```
id
210100E08B2520D4
```

lshost

The **lshost** command generates a list with concise information about all the hosts visible to the cluster and detailed information about a single host.

Syntax

```
svcinfo lshost [-filtervalue attrib=value] [-nohdr]
                [-delim delimiter] [-filtervalue?] [object_id | object_name]
```

Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- When using a wildcard character, you must enclose the filter entry within double quotation marks (" "), as follows:

```
svcinfo lshost -filtervalue "name=md*"
```

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | *object_name*

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id* | *object_name* parameter, the concise view of all objects matching the filtering requirements that is specified by the **-filtervalue** parameter are displayed.

-filtervalue?

(Optional) Specifies that you want your report to display any or all of the list of valid filter attributes. The valid filter attributes for the **svcinfo lshost** command are:

- host_name
- host_id
- port_count
- name
- id
- iogrp_count

Description

This command returns a concise list or a detailed view, of hosts visible to the cluster.

The following list provides the different states for a host:

offline

| The host is offline if one or more I/O groups with VDisk (volume) mappings do not have a login for the specified WWPN.

| **Note:** iSCSI hosts are always displayed in an offline state, even when they are functioning correctly. This is a current limitation for iSCSI hosts.

degraded

The host is degraded if one or more nodes with VDisk mappings do not have a login for the specified WWPN.

inactive

The host is inactive if all the nodes with VDisk mappings have a login for the specified WWPN, however, no nodes have seen any SCSI commands from the WWPN in the last 5 minutes.

active The host is active if all the nodes with VDisk mappings have a login for the specified WWPN, however, at least one node has seen SCSI commands from the WWPN in the last 5 minutes.

If a host does not have any VDisk mappings, then it is either reported as offline or inactive.

Note: The `svcinfo lshost` command presents a list of host HBA ports that are logged in to nodes. However, there are situations when the information presented can include host HBA ports that are no longer logged in or even part of the SAN fabric. For example: A host HBA port is unplugged from a switch, but `svcinfo lshost` still shows the WWPN logged in to all nodes. If this occurs, the incorrect entry is removed when another device is plugged in to the same switch port that previously contained the removed host HBA port.

A concise invocation example

```
svcinfo lshost
```

The concise resulting output

id	name	port_count	iogrp_count
0	mchost20	1	4
1	mchost30	1	4
2	mchost200	1	4
3	mchost40	1	4
4	mchost240	1	4
5	mchost170	1	4
6	mchost120	1	4
7	mchost60	1	4
8	mchost180	1	4
9	mchost13	2	4

A detailed invocation example

```
svcinfo lshost mchost13
```

The detailed resulting output

```
id 9
name mchost13
port_count 2
type generic
mask 1111
iogrp_count 4
iscsiname mc13host
node_logged_in_count 0
state offline
iscsiname mchost13
node_logged_in_count 0
state offline
```

lshostiogrp

The **lshostiogrp** command displays a list of all the I/O groups that are associated with a specified host.

Syntax

```
➤➤ svcinfo -- lshostiogrp -- [-nohdr] [-delim delimiter] [host_id | host_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

host_id | *host_name*

(Required) The name or ID of the host for which the list of I/O groups is required.

Description

This command displays a list of all the I/O groups that are mapped to the specified host.

An invocation example

```
svcinfo lshostiogrp -delim : hostone
```


The resulting output

```
id:name
0:io_grp0
1:io_grp1
```

Ishostvdiskmap

The **Ishostvdiskmap** command displays a list of VDIs that are mapped to a given host. These are the VDIs that are recognized by the specified host.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

host_id | host_name

(Optional) Specifies the host in terms of its ID or name. The command displays a list of all the virtual disks that are mapped to the specified host and the SCSI ID by which they are mapped. If neither a host ID or name are entered, the command displays a list of all recognized host-to-VDisk mappings.

Description

This command displays a list of VDisk IDs and names. These are the VDIs that have been mapped to the specified host; that is, they are visible to the specified host. The SCSI LUN ID is also displayed. This SCSI LUN ID is the ID by which the VDisk is recognized by the host.

Each VDisk that is exported by the cluster is assigned a unique vpath number. This number identifies the VDisk and determines which VDisk corresponds to the volume that the hosts recognize. This procedure can only be performed using the command-line interface.

For the specified volume, find the vpath serial number by issuing the **datapath query device** command. Find the host that is defined to the cluster that corresponds with the host that you are working with.

1. The WWPNs are an attribute of the HBA. You can find these by looking at the device definitions stored by your operating system. For example, on AIX they are in the ODM, in Windows they are in the Device Manager details for the given HBA.
2. Verify which host is defined to the cluster that these ports belong to. The ports are stored as part of the detailed view, so you must list each host in turn by issuing the following command:

```
svcinfo lshost host_name | host_id
```

where *host_name* | *host_id* is the name or ID of the host. Check for matching WWPNs.

Note: Name your hosts accordingly. For example, if the actual host is called **orange**, also name the host that is defined to the cluster **orange**.

When you have the *hostname* defined to the cluster and the *vpath serial number*, issue the following command:

```
svcinfolshostvdiskmap hostname
```

where *hostname* is the name of the host. A list is displayed. Look for the VDisk UID that matches the *vpath serial number* and record the VDisk name or ID.

An invocation example

```
svcinfolshostvdiskmap -delim : 2
```

The resulting output

```
id:name:SCSI_id:vdisk_id:vdisk_name:vdisk_UID
2:host2:0:10:vdisk10:6005076801958001500000000000000A
2:host2:1:11:vdisk11:6005076801958001500000000000000B
2:host2:2:12:vdisk12:6005076801958001500000000000000C
2:host2:3:13:vdisk13:6005076801958001500000000000000D
2:host2:4:14:vdisk14:6005076801958001500000000000000E
```

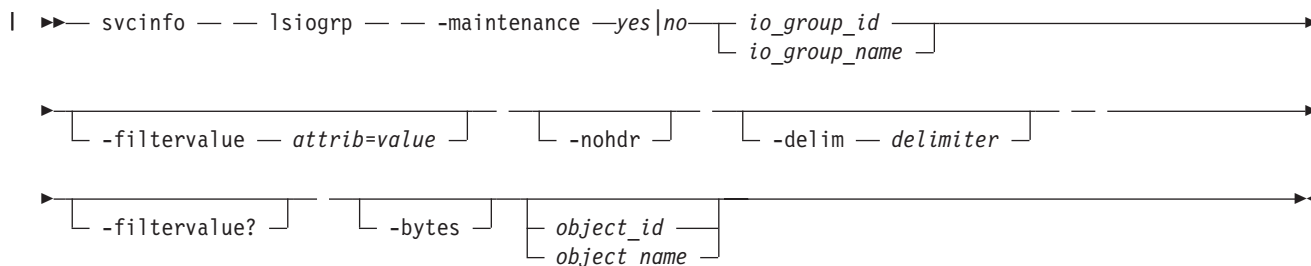
Isiogrp

The **lsiogrp** command returns a concise list or a detailed view of I/O groups visible to the cluster.

The list report style can be used to obtain the following two styles of report:

- A list containing concise information about all the I/O groups that are visible to the cluster. Each entry in the list corresponds to a single I/O group.
- The detailed information about a single I/O group.

Syntax



Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcard characters with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*), which must be the first or last character in the string.
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:

```
svcinfolsiogrp -filtervalue "name=md*"
```

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-filtervalue?

(Optional) Displays the valid filter attributes for the **svcinfolsiogrp** command:

- | • **HWS_name**
- | • **HWS_unique_id**
- | • **node_count**
- | • **name**
- | • **id**
- | • **host_count**

-bytes

(Optional) Displays all capacities as bytes.

object_id | *object_name*

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id* | *object_name* parameter, the concise view of all objects matching the filtering requirements that is specified by the **-filtervalue** parameter are displayed.

| **-maintenance** *yes* | *no*

(Optional) Specifies whether the I/O group should be in maintenance mode. The I/O group should be placed in maintenance mode while carrying out service procedures on storage enclosures. Once you enter maintenance mode, it continues until either:

- | • It is explicitly cleared, OR
- | • 30 minutes elapse

Description

This command returns a concise list or a detailed view of I/O groups visible to the cluster.

| You can display the following information for this command:

| **raid_total_memory**

| Total bitmap space available for RAID arrays (in MB with 1 decimal place).

| **raid_free_memory**

| Bitmap space available for creating new RAID arrays (in MB with 1 decimal place).

A concise invocation example

```
svcinfolsiogrp -delim :
```

The concise resulting output

```
id:name:node_count:vdisk_count:host_count
0:io_grp0:1:0:0
1:io_grp1:0:0:0
2:io_grp2:0:0:0
3:io_grp3:0:0:0
4:recovery_io_grp:0:0:0
```

A detailed invocation example

```
svcinflsiogrp -delim : 0
```

The detailed resulting output

```
id:0
name:io_grp0
node_count:1
vdisk_count:51
host_count:0
flash_copy_total_memory:3.0MB
flash_copy_free_memory:1.0MB
remote_copy_total_memory:6.5MB
remote_copy_free_memory:2.8MB
mirroring_total_memory:1.0MB
mirroring_free_memory:0.3MB
maintenance:no
```

Isiogrp host

The **Isiogrp host** command displays a list of the hosts that are mapped to a specified I/O group.

Syntax

```
svcinflsiogrp host [-nohdr] [-delim delimiter] [iogrp_id | iogrp_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

iogrp_id | iogrp name

(Required) The ID or name of the I/O group for which a list of all mapped hosts is required.

Description

The **lsiogrphost** command displays a list of hosts that are mapped to a specified I/O group.

An invocation example

```
svcinflsiogrphost -delim : 0
```

The resulting output

```
id:name
0:hostzero
1:hostone
```

lsiogrpcandidate

Use the **lsiogrpcandidate** command to list the I/O groups that can have nodes added to them.

Syntax

```
▶▶ svcinfo — — lsiogrpcandidate — — [ -nohdr ] [ -delim — delimiter ] ▶▶
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays a list of I/O groups to which nodes can be added. Only the I/O group IDs are displayed.

An invocation example

```
svcinflsiogrpcandidate
```

The resulting output

```
id
0
1
2
3
4
```

Isiostatsdumps (Deprecated)

| **Attention:** The `svcinfolisiostatsdumps` command is deprecated. Use the `svcinfolsdumps` command to display a list of files in a particular dumps directory.

Isiotracedumps (Deprecated)

| **Attention:** The `svcinfolsiotracedumps` command is deprecated. Use the `svcinfolsdumps` command to display a list of files in a particular dumps directory.

Isiscsiauth

The `Isiscsiauth` command lists the Challenge Handshake Authentication Protocol (CHAP) secret configured for authenticating an entity to the SAN Volume Controller cluster.

Syntax

```
svcinfolsiiscsiauth [-nohdr] [-delim delimiter] [-filtervalue attribute=value] [-filtervalue?]
```

Parameters

`-nohdr`

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The `-nohdr` parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

`-delim delimiter`

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The `-delim` parameter overrides this behavior. Valid input for the `-delim` parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

`-filtervalue attribute=value`

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed. If a capacity is specified, the units must also be included.

Note: Some filters allow the asterisk character (`*`) when you enter the command. The following rules apply to the use of wildcard characters with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (`*`).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, you must enclose the filter entry within double quotation marks (`"`), as follows:

```
svcinfolsiiscsiauth -filtervalue "name=md*"
```


-filtervalue?

(Optional) displays a list of filters that can be applied against this view. The following filter attributes are valid for the **lisiscsiauth** command:

- **type**
- **id**
- **name**
- **iscsi_auth_method**
- **iscsi_chap_secret**

Description

This command lists the CHAP secret configured for authenticating an entity to the SAN Volume Controller cluster. The command also displays the configured iSCSI authentication method. The `iscsi_auth_method` field can have values of `none` or `chap`.

When you create an iSCSI host using the **svctask mkhost** command with the **iscsiname** parameter, the host is initially configured with the authentication method as `none`, and no CHAP secret is set. To set a CHAP secret for authenticating the iSCSI host with the SAN Volume Controller cluster, use the **svctask chhost** command with the **chapsecret** parameter.

A invocation example

```
svcinfolisiscsiauth
```

The resulting output

type	id	name	iscsi_auth_method	iscsi_chap_secret
host	0	mhost20	none	
host	1	mhost30	none	
host	2	mhost200	none	
host	3	mhost40	none	
host	4	mhost240	none	
host	5	mhost170	none	
host	6	mhost120	none	
host	7	mhost60	none	
host	8	mhost180	none	
host	9	mhost13	none	
host	10	newhost	none	

lslicense

The **lslicense** command displays current license settings for cluster features.

Syntax

```
svcinfolicense -nohdr -delim delimiter
```

Parameters

-nohdr

(Optional) Suppresses the display of these headings. By default, headings are displayed for each column of data (in a concise style view providing general information about objects of a particular type) and for each item of data (in a detailed style view providing much more information about a specific object of a particular type).

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :`, a colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

| The **lslicense** command displays license settings for cluster features, including remote copy and virtualization settings. SAN Volume Controller also includes FlashCopy settings. The displayed output for SAN Volume Controller lists capacity values in terabytes (TB) and feature enablement. The displayed output for Storwize V7000 lists enclosure license values.

Use the **chlicense** command to change the feature license settings. Because the feature license settings are entered when the cluster is first created, you must only update the settings if you have changed your license.

An invocation example

```
svcinfolicense
```

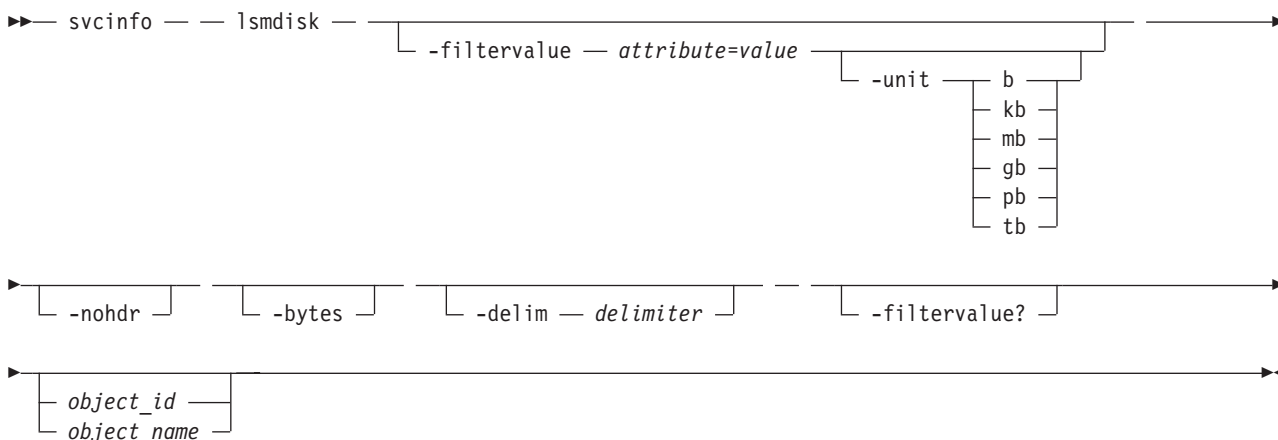
The resulting output

```
| used_flash 0.00  
| used_remote 0.00  
| used_virtualization 0.00  
| license_flash 0  
| license_remote 20  
| license_virtualization 30  
| license_physical_disks 0  
| license_physical_flash off  
| license_physical_remote off
```

lsmdisk

The **lsmdisk** command returns a concise list or a detailed view of managed disks (MDisks) visible to the cluster. It can also list detailed information about a single MDisk.

Syntax



Parameters

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filter attributes matching the specified values; see **-filtervalue?** for the supported attributes. Only objects with a value that matches the filter attribute value are returned. If **capacity** is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when entering the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- When using a wildcard character, you must enclose the filter entry within double quotation marks (""), as follows:

```
svcinfolsmdisk -filtervalue "name=md*"
```

-filtervalue?

(Optional) Includes all of the valid filter attributes in the report. The following filter attributes are valid for the **svcinfolsmdisk** command:

- | • id
- | • name
- | • status
- | • mode
- | • mdisk_grp_id
- | • mdisk_grp_name
- | • capacity
- | • quorum_index
- | • block_size
- | • controller_name
- | • ctrl_WWNN
- | • controller_id
- | • path_count
- | • ctrl_LUN_#
- | • UID
- | • preferred_WWPN
- | • active_WWPN
- | • tier

Any parameters specified with the **-filtervalue?** parameter are ignored.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-bytes

(Optional) Specifies that you want the report to display all capacities as bytes. Capacity values displayed in units other than bytes might be rounded. When filtering on capacity, use a unit of bytes, **-unit b**, for exact filtering.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space separated. The width of each

column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | object_name

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id | object_name* parameter, the concise view displays all objects matching the filtering requirements that are specified by the **-filtervalue** parameter.

Description

This command returns a concise list or a detailed view of MDisks visible to the cluster. Table 30 provides the potential output for MDisks.

Table 30. MDisk output

Attribute	Values
status	<ul style="list-style-type: none"> • online • offline • excluded • degraded_paths • degraded_ports • degraded (applies only to internal MDisks)
mode	unmanaged, managed, image, array
quorum_index	0, 1, 2, or blank if the MDisk is not being used as a quorum disk
block_size	512, 524 bytes in each block of storage
ctrl_type	4, 6, where 6 is a solid-state drive (SSD) attached inside a node and 4 is any other device
tier	<p>The tier this MDisk has been assigned to by auto-detection (for internal arrays) or by the user:</p> <ul style="list-style-type: none"> • generic_ssd • generic_hdd (the default value for newly discovered or external MDisk) <p>Note: You can change this value using the <code>chmdisk</code> command.</p>
raid_status	<ul style="list-style-type: none"> • offline - the array is offline on all nodes • degraded - the array has deconfigured or offline members; the array is not fully redundant • syncing - array members are all online, the array is syncing parity or mirrors to achieve redundancy • initting - array members are all online, the array is initializing; the array is fully redundant • online - array members are all online, and the array is fully redundant
raid_level	The RAID level of the array (RAID0, RAID1, RAID5, RAID6, RAID10).
redundancy	The number of how many member disks can fail before the array fails.
strip_size	The strip size of the array (in KB).
spare_goal	The number of spares that the array members should be protected by.
spare_protection_min	The minimum number of spares that an array member is protected by.

Table 30. MDisk output (continued)

Attribute	Values
balanced	Describes if the array is balanced to its spare goals: <ul style="list-style-type: none"> • exact: all populated members have exact capability match, exact location match • yes: all populated members have at least exact capability match, exact chain, or different enclosure or slot • no: anything else

Note: The automatic discovery performed by the cluster does not write anything to an unmanaged MDisk. It is only when you add an MDisk to an MDisk group (storage pool), or use an MDisk to create an image mode VDisk (volume), that the system uses the storage.

To see which MDisks are available, issue the **svctask detectmdisk** command to manually rescan the fibre-channel network for any new MDisks. Issue the **svcinfo lsmdiskcandidate** command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group (storage pool).

Notes:

1. The MDisk *path_count* value is the number of nodes through which the MDisk is currently accessible.
2. The MDisk *max_path_count* value is the highest value that the *path-count* has reached since the MDisk was initially discovered or since the MDisk was last accessible through all nodes in the cluster.

The following define the status fields:

Online

The MDisk is online and available.

Degraded

(Internal MDisks only) The array has members that are degraded, or the *raid_status* is degraded.

Degraded ports

There are one or more MDisk port errors.

Degraded paths

One or more paths to the MDisk have been lost; the MDisk is not online to every node in the cluster.

Offline

All paths to the MDisk are lost.

Excluded

The MDisk is excluded from use by the cluster; the MDisk port error count exceeded the threshold.

A concise invocation example

```
svcinfo lsmdisk -delim :
```

The concise resulting output

```
| id:name:status:mode:mdisk_grp_id:mdisk_grp_name:capacity:ctrl_LUN_#:controller_name:UID:tier
0:mdisk0:online:unmanaged:::68.4GB:0000000000000000:controller0:
20000004cf2422aa00000000000000000000000000000000000000000000000000000000:generic_hdd
1:mdisk1:online:unmanaged:::68.4GB:0000000000000000:
controller1:20000004cf1fd19d00000000000000000000000000000000000000000000000000000000:generic_hdd
2:mdisk2:online:unmanaged:::68.4GB:0000000000000000:controller2:
20000004cf2425310000000000000000000000000000000000000000000000000000000:generic_hdd
```

A detailed invocation example

```
| svcinfo lsmdisk mdisk1
```

The detailed resulting output

```
| id:1
| name:mdisk1
| status:online
| mode:array
| mdisk_grp_id:0
| mdisk_grp_name:mdgp0
| capacity:136.0GB
| quorum_index:
| block_size:512
| controller_name:controller1
| ctrl_type:4
| ctrl_WWNN:200400A0B80F0702
| controller_id:1
| path_count:2
| max_path_count:2
| ctrl_LUN #:0000000000000002
| UID:600a0b80000f07020000005c45ff8a7c00000000000000000000000000000000
| preferred_WWPN:200400A0B80F0703
| active_WWPN:200400A0B80F0703
| node_id:
| node_name:
| location:
| fast_write_state:empty
| raid_status:
| raid_level:
| redundancy:
| strip_size:
| spare_goal:
| spare_protection_min:
| balanced:
| tier:generic_hdd
```

A detailed invocation example

```
| svcinfo lsarray mdisk3
```

The resulting output

```
| mdisk_id:3
| mdisk_name:mdisk3
| status:online
| mode:array
| mdisk_grp_id:0
| mdisk_grp_name:mdiskgrp0
| capacity:68.4GB
| quorum_index:
| block_size:
| controller_name:
| ctrl_type:
| ctrl_WWNN:
| controller_id:
| path_count:
| max_path_count:
| ctrl_LUN #:
| UID:
| preferred_WWPN:
| active_WWPN:
| node_id:
| node_name:
| location:
| fast_write_state:empty
| raid_status:online
| raid_level:raid0
| redundancy:0
| strip_size:256
```

```
| spare_goal:2
| spare_protection_min:2
| balanced:yes
| tier:generic_ssd
```

Ismdiskdumps (Deprecated)

| **Attention:** The `svcinfolismdiskdumps` command is deprecated. Use the `svcinfolsdumps` command to display a list of files in a particular dumps directory.

Ismdisklba

The `Ismdisklba` command lists the MDisk and logical block address (LBA) for the specified VDisk LBA.

Syntax

```
►►► svcinfo — — Ismdisklba — — -vdisklba — vdisklba —————►
                                   └ -copy — id ┘
└ -delim — delimiter ┘ — -nohdr — -vdisk — vdisk_id —————►
                                   └ vdisk_name ┘
```

Parameters

-vdisklba *vdisklba*

(Required) Specifies the 64-bit hexadecimal logical block address (LBA) on the VDisk.

-copy *id*

(Optional) Specifies the VDisk copy ID to list the MDisk and LBA for. If this parameter is not specified, the command lists MDisks and LBAs for all VDisk copies.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

vdisk_id | *vdisk_name*

(Required) Specifies the VDisk name or ID.

Description

The `Ismdisklba` command returns the logical block address (LBA) of the MDisk that is associated with the VDisk LBA. For mirrored VDIsks, the command lists the MDisk LBA for both the primary and the copy.

If applicable, the command also lists the range of LBAs on both the VDisk and MDisk that are mapped in the same extent, or for space-efficient disks, in the same grain. If a space-efficient VDisk is offline and the specified LBA is not allocated, the command displays the VDisk LBA range only.

Table 31 summarizes the data that can be returned with this command.

Table 31. *lsmdisklba* command output

Field	Fully allocated, single copy VDisk	LBA not allocated on space-efficient VDisk	Mirrored VDisk with one normal copy and one offline space-efficient copy	
			Normal copy	Space-efficient copy
<code>copy_id</code>	yes	yes	yes	yes
<code>mdisk_id</code>	yes	no	yes	no
<code>mdisk_name</code>	yes	no	yes	no
<code>type</code>	allocated	unallocated	allocated	offline
<code>mdisk_lba</code>	yes	no	yes	no
<code>mdisk_start</code>	yes	no	yes	no
<code>mdisk_end</code>	yes	no	yes	no
<code>vdisk_start</code>	yes	yes	yes	yes
<code>vdisk_end</code>	yes	yes	yes	yes

An invocation example

```
svcinfolsmdisklba -vdisk 0 -vdisklba 0x0
```

The resulting output

```
copy_id 0
mdisk_id 1
mdisk_name mdisk1
type allocated
mdisk_lba 0x00090000
mdisk_start 0x00090000
mdisk_end 0x000907FF
vdisk_start 0x00000000
vdisk_end 0x000007FF
```

lsmdiskcandidate

The `lsmdiskcandidate` command lists all of the unmanaged MDisks by MDisk ID.

Syntax

```
svcinfolsmdiskcandidate [-nohdr] [-delim delimiter]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The `-nohdr` parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays a list of MDisks that are unmanaged. Only the MDisk IDs are displayed.

When back-end controllers are added to the fibre-channel SAN and are included in the same switch zone as a cluster, the cluster automatically detects the back-end controller to determine which storage is presented to the node. The SCSI logical units that are presented by the back-end controller are displayed as unmanaged MDisks. However, if the configuration of the back-end controller is modified after this has occurred, the cluster might be unaware of these configuration changes. You can then request that the cluster rescan the fibre-channel SAN to update the list of unmanaged MDisks.

Note: The automatic detection performed by the cluster does not write anything to a unmanaged MDisk. It is only when you instruct the cluster to add an MDisk to a managed disk group or use a MDisk to create an image mode virtual disk that the storage is actually used.

Check to see which MDisks are available by issuing the `svctask detectmdisk` command to manually scan the fibre-channel network for any MDisks. Issue the `svcinfo lsmdiskcandidate` command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group. Alternatively, you can issue the `svcinfo lsmdisk` command to view all of the MDisks.

An invocation example

```
svcinfo lsmdiskcandidate
```

The resulting output

```
id
5
6
7
8
9
10
11
12
13
14
```

lsmdiskextent

The `lsmdiskextent` command displays the extent allocation between managed disks and virtual disks. The output lists a VDisk ID, VDisk copy ID, and the number of extents.

Syntax

```
▶▶ svcinfo — — lsmdiskextent — [ -nohdr ] [ -delim — delimiter ] —————▶
```



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

mdisk_name | *mdisk_id*

(Required) Specifies the specific object ID or name of the given type.

Description

The command displays a list, in which each entry contains a VDisk ID, VDisk copy ID, and the number of extents. These VDisk copies are using extents on the specified MDisk. The number of extents being used on each MDisk is also shown.

Every VDisk copy is constructed from one or more MDisks. At times, you might have to determine the relationship between the two objects. The following procedure allows you to determine the relationships.

To determine the relationship between VDisk copies and MDisks, issue the following command for each VDisk copy:

```
svcinfolsvdiskmember vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the VDisk copy. This displays a list of IDs that correspond to the MDisks that make up the VDisk copy.

To determine the relationship between VDisk copies and MDisks and the number of extents that are provided by each MDisk, you must use the command-line interface. For each VDisk copy, issue the following command:

```
svcinfolsvdiskextent vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the VDisk copy. This displays a table of MDisk IDs and the corresponding number of extents that each MDisk is providing as storage for the given VDisk copy.

To determine the relationship between MDisks and VDisk copies, issue the following command for each MDisk:

```
svcinfolsmdiskmember mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This displays a list of IDs that correspond to the VDisk copies that are using this MDisk.

To determine the relationship between MDisks and VDisk copies and the number of extents that are used by each VDisk copy, you must use the command-line interface. For each MDisk, issue the following command:

```
svcinfolsmdiskextent mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This command displays a table of VDisk copy IDs and the corresponding number of extents that are being used by each VDisk copy.

An invocation example

```
svcinfolsmdiskextent -delim : mdisk0
```


The resulting output

```
id:number_of_extents:copy_id
1:1:1
```

lsmdiskgrp

- | The **lsmdiskgrp** command returns a concise list or a detailed view of MDisk groups (storage pools) that are visible to the cluster.

Syntax



Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- When using a wildcard, you must enclose the filter entry within double quotation marks (""), as follows:

```
svcinfolsmdiskgrp -filtervalue "name=md*"
```

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-bytes

(Optional) Specifies that you want the report to display all capacities as bytes.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | *object_name*

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id* | *object_name* parameter, the concise view of all objects matching the filtering requirements specified by the **-filtervalue** parameter are displayed.

-filtervalue?

Displays a list of valid filter attributes. The valid filters for the `svcinfo lsmdiskgrp` command are:

- name
- storage_pool_id
- mdisk_count
- vdisk_count
- extent_size
- status
- id
- | • easy_tier
- | • easy_tier_status

Description

This command returns a concise list or a detailed view of MDisk groups visible to the cluster.

Command output includes values for the following attributes:

status The state of the MDisk with the highest-priority status in the group, excluding image mode MDisks.

VDisk_count

The number of VDisk (volume) copies that are in the MDisk group (storage pool).

capacity

The total amount of MDisk storage that is assigned to the MDisk group.

extent_size

The size of the extents for this group: 16, 32, 64, 128, 256, 512, 1024, 2048, 4096, or 8192 (MB).

free_capacity

The amount of MDisk storage that is assigned to the MDisk group that is unused. MDisk storage can be used for cluster quorum data and VDIsks.

real_capacity

The amount of MDisk storage that is assigned to the MDisk group that is assigned to VDIsks.

virtual_capacity

The total virtual size of all the VDisk copies that are associated with the MDisk group. This is the same as the `real_capacity` value unless you have configured space-efficient VDisk copies in this MDisk group.

used_capacity

The total used size of all the VDisk copies that are associated with the MDisk group. This is the same as the real_capacity value unless you have configured space-efficient VDisk copies in this MDisk group.

overalllocation

Expressed as a percentage, the ratio of the virtual_capacity value to the capacity. An MDisk group overalllocation of over 100 is only possible if you have configured space-efficient VDisk copies.

warning

This field is a percentage. A warning is generated when the amount of space in the MDisk group that has been assigned exceeds this level.

| easy_tier

This value is set by the user and determines whether Easy Tier is permitted to manage the pool.

| Note:

- | 1. If easy_tier is *on*, then easy_tier_status is *active*
- | 2. If easy_tier is *off*, then easy_tier_status is *inactive*
- | 3. If easy_tier is *auto*, then the value of easy_tier_status is determined by the number of tiers an MDisk group has.

| easy_tier_status

Whether the Easy Tier functions are active on an MDisk group:

- | • active
- | • inactive

| tier

Which tier information is being reported:

- | • generic_ssd
- | • generic_hdd

| tier_mdisk_count

The number of MDisks in the tier.

| tier_capacity

The total MDisk capacity assigned to the VDisk in the tier.

Note: For space-efficient copies, the capacity by tier will be the real capacity.

| tier_free_capacity

The unused amount of MDisk storage in the tier.

The following define the status fields, from lowest to highest priority:

Online

The MDisk group is online and available.

Offline

All paths to the MDisk group are lost.

A concise invocation example

```
svcinfolsmdiskgrp -delim :
```

The concise resulting output

```
| id:name:status:mdisk_count:vdisk_count:capacity:extent_size:free_capacity:  
| virtual_capacity:used_capacity:real_capacity:overalllocation:warning:easy_tier:easy_tier_status  
| 0:mdiskgrp0:degraded:4:0:34.2GB:16:34.2GB:0:0:0:0:auto:inactive  
| 1:mdiskgrp1:online:4:6:200GB:16:100GB:400GB:75GB:100GB:200:80:on:active
```

| An detailed invocation example for an MDisk group with one tier

```
| svcinfo lsmdiskgrp -delim : mdiskgrp1
```

| **The resulting output**

```
| id:1  
| name:mdiskgrp1  
| status:online  
| mdisk_count:4  
| vdisk_count:6  
| capacity:200GB  
| extent_size:16  
| free_capacity:100GB  
| virtual_capacity:400.00GB  
| used_capacity:75.00GB  
| real_capacity:100.00GB  
| overallocation:200  
| warning:80  
| easy_tier:on  
| easy_tier_status:active  
| tier:generic_ssd  
| tier_mdisk_count:0  
| tier_capacity: 0.00MB  
| tier_free_capacity: 0.00MB  
| tier:generic_hdd  
| tier_mdisk_count:4  
| tier_capacity:200.00GB  
| tier_free_capacity:100.00GB
```

| **An detailed invocation example for an MDisk group with two tiers**

```
| svcinfo lsmdiskgrp -delim : mdiskgrp2
```


| **The resulting output**

```
| id:2  
| name:mdiskgrp2  
| status:online  
| mdisk_count:8  
| vdisk_count:6  
| capacity:200GB  
| extent_size:16  
| free_capacity:100GB  
| virtual_capacity:400.00GB  
| used_capacity:75.00GB  
| real_capacity:100.00GB  
| overallocation:200  
| warning:80  
| easy_tier:auto  
| easy_tier_status:active  
| tier:generic_ssd  
| tier_mdisk_count:2  
| tier_capacity:20.00GB  
| tier_free_capacity: 0.00MB  
| tier:generic_hdd  
| tier_mdisk_count:6  
| tier_capacity:180.00GB  
| tier_free_capacity:100.00GB
```

lsmdiskmember

The **lsmdiskmember** command displays a list of VDIsks that are using extents on the specified MDisk. That is, the virtual disks are using extents on the managed disk that are specified by the MDisk ID.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

mdisk_id | *mdisk_name*

(Required) Specifies the ID or name of the MDisk for which you want a list of VDisks that use extents of that MDisk.

Description

This command displays a list of virtual disks that use extents on the managed disk that are specified by the ID. The list displays members of the respective object and is independent of the state of the individual members; that is, if they are in offline state, they are still displayed.

Every VDisk is constructed from one or more MDisks. To determine the relationship between VDisk copies and MDisks, issue the following command:

```
svcinfo lsvdiskmember vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the VDisk copy. This displays a list of IDs that correspond to the MDisks that make up the VDisk copy.

To determine the relationship between VDisk copies and MDisks and the number of extents that are provided by each MDisk, you must use the command-line interface. For each VDisk copy, issue the following command:

```
svcinfo lsvdiskextent vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the VDisk copy. This command displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the VDisk copy.

To determine the relationship between MDisks and VDisk copies, issue the following command:

```
svcinfo lsmdiskmember mdisk_id | mdisk_name
```

where *mdisk_id* | *mdisk_name* is the name or ID of the MDisk. This command displays a list of IDs that correspond to the VDisk copies that are using this MDisk.

To determine the relationship between MDisks and VDisk copies and the number of extents that are used by each VDisk copy, you must use the command-line interface. For each MDisk *mdisk_id* | *mdisk_name*, issue the following command:

```
svcinfolsmdiskextent mdisk_id | mdisk_name
```

where *mdisk_id* | *mdisk_name* is the name or ID of the MDisk. This command displays a table of VDisk copy IDs and the corresponding number of extents that are being used by each VDisk copy.

An invocation example

```
svcinfolsmdiskmember -delim : 1
```

The resulting output

```
id:copy_id
0:0
1:0
2:0
3:0
4:0
5:0
6:0
```

lsmigrate

The **lsmigrate** command displays the progress of all current data migration operations.

Syntax

```
➤— svcinfolsmigrate — [ -nohdr ] [ -delim delimiter ] ➤
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

If you use multiple threads to migrate data, the progress will increment when all threads have completed the migration of an extent. For large extent sizes with many threads, this can result in quite large increments in the percentage progress.

Description

This command displays information of all the migrations that are currently in progress.

Note: Only user-initiated migrations that are reported using this command. Easy Tier migrations are not included in the output.

An invocation example

```
svcinfo lsmigrate -delim :
```

The resulting output

```
migrate_type:MDisk_Group_Migration
progress:96
migrate_source_vdisk_index:33
migrate_target_mdisk_grp:4
max_thread_count:4
migrate_source_vdisk_copy_id:1
```


lsnode

The **lsnode** command returns a concise list or a detailed view of nodes that are part of the cluster.

The list report style can be used to obtain two styles of report:

- A list containing concise information about all the nodes on a cluster. (Each entry in the list corresponds to a single node.)
- The detailed information about a single node.

Syntax



Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- When using a wildcard, you must enclose the filter entry within double quotation marks (""), as follows:

```
svcinfo lsnode -filtervalue "name=md*"
```

-filtervalue?

Displays a list of valid filter attributes for the **-filtervalue** *attribute=value* parameter. The valid filters for the **svcinfo lsnode** command are:

- node_name
- id

- status
- IO_group_name
- IO_group_id
- name
- hardware
- | • service_IP_address
- | • UPS_serial_number
- | • WWNN
- | • partner_node_id
- | • partner_node_name
- | • config_node
- | • UPS_unique_id

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | *object_name*

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id* | *object_name* parameter, the concise view of all objects matching the filtering requirements that is specified by the **-filtervalue** parameter are displayed.

Description

SAN Volume Controller: This command returns a concise list or a detailed view of nodes that are part of the cluster. Table 32 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 32. *svcinfo lsnode* attribute values

Attribute	Value
status	offline flushing pending online adding deleting
config_node	no yes
port_status	active inactive not installed
hardware	8F2 8F4 8G4 CF8 8A4 other
UPS_serial_number	The serial number of the UPS.
status	The status of the node.
UPS_unique_id	The unique ID of the UPS.
panel_name	Unique identifier for the nodes.

Table 32. svcinfo lsnode attribute values (continued)

Attribute	Value
enclosure_id	Blank.
canister_id	Blank.
enclosure_serial_number	Blank.

Storwize® V7000: This command returns a concise list or a detailed view of nodes that are part of the cluster. Table 33 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 33. svcinfo lsnode attribute values

Attribute	Value
status	offline flushing pending online adding deleting
config_node	no yes
port_status	active inactive not installed
hardware	8F2 8F4 8G4 CF8 8A4 other
UPS_serial_number	Blank.
status	The status of the node.
UPS_unique_id	Blank.
panel_name	A compound of enclosure ID and canister ID, separated with a hyphen.
enclosure_id	The enclosure that the node is in, shown on the front panel.
canister_id	The canister that the node is in.
enclosure_serial_number	The serial number of the enclosure.

A concise invocation example

```
svcinfo lsnode -delim
```

The resulting SAN Volume Controller output:

```
| IBM_2145:ldcluster-19:admin>svcinfo lsnode -delim ,
| id,name,UPS_serial_number,WNN,status,IO_group_id,IO_group_name,config_node,UPS_unique_id,hardware,
| iscsi_name,
| iscsi_alias,panel_name,enclosure_id,canister_id,enclosure_serial_number1,hlcn114289,
| 10004BC018,5005076801002978,
| online,0,io_grp0,no,20400001124C0048,8A4,iqn.1986-03.com.ibm:2145.ldcluster-19.hlcn114289,,
| 114289,,
| 2,hlcn114253,10004BC023,5005076801002822,online,0,io_grp0,yes,20400001124C0083,8A4,
| iqn.1986-03.com.ibm:2145.ldcluster-19.hlcn114253,,114253,,5,hdn116511,1000871087,
| 5005076801005CCE,online,
| 1,io_grp1,no,2040000207040207,CF8,iqn.1986-03.com.ibm:2145.ldcluster-19.hdn116511,,116511,,
| 4,hdn116520,100062L106,5005076801005FD0,online,1,io_grp1,no,2040000182701006,CF8,
| iqn.1986-03.com.ibm:2145.ldcluster-19.hdn116520,,116520,,
```

The resulting Storwize V7000 output:

```
| id,name,UPS_serial_number,WNN,status,IO_group_id,IO_group_name,config_node,UPS_unique_id,hardware,
| iscsi_name,iscsi_alias,
| panel_name,enclosure_id,canister_id,enclosure_serial_number1,node1,,50050768010010FE,online,0,
| io_grp0,yes,
| 50050768010010FE,100,iqn.1986-03.com.ibm:2145.tbcluster-12.node1,,01-1,1,1,00001A2
| 2,node2,,5005076801000FFE,online,0,io_grp0,no,5005076801000FFE,100,iqn.1986-03.com.ibm:
| 2145.tbcluster-12.node2,,
| 01-2,1,2,00001A2
```

| A detailed invocation example

```
svcinfolsnode -delim , 1
```

```
| svcinfolsnode 1
```

The resulting SAN Volume Controller output

```
id,1
name,h1cn114289
UPS_serial_number,10004BC018
WWNN,5005076801002978
status,online
IO_group_id,0
IO_group_name,io_grp0
partner_node_id,2
partner_node_name,h1cn114253
config_node,no
UPS_unique_id,20400001124C0048
port_id,5005076801402978
port_status,active
port_speed,4Gb
port_id,5005076801302978
port_status,active
port_speed,4Gb
port_id,5005076801102978
port_status,active
port_speed,4Gb
port_id,5005076801202978
port_status,active
port_speed,4Gb
hardware,8A4
iscsi_name,iqn.1986-03.com.ibm:2145.lcluster-19.h1cn114289
iscsi_alias,
failover_active,no
failover_name,h1cn114253
failover_iscsi_name,iqn.1986-03.com.ibm:2145.lcluster-19.h1cn114253
failover_iscsi_alias,
panel_name,114289
enclosure_id,
canister_id,
enclosure_serial_number,
service_IP_address,9.180.29.52
service_gateway,9.180.28.1
service_subnet_mask,255.255.254.0
service_IP_address_6,
service_gateway_6,
service_prefix_6,
```

| The Storwize V7000 resulting output:

```
| id 1
| name node1
| UPS_serial_number
| WWNN 50050768010010FE
| status online
| IO_group_id 0
| IO_group_name io_grp0
| partner_node_id 2
| partner_node_name node2
| config_node yes
| UPS_unique_id 50050768010010FE
| port_id 50050768011010FE
| port_status active
| port_speed 4Gb
| port_id 50050768012010FE
| port_status active
| port_speed 4Gb
```

```

| port_id 50050768013010FE
| port_status active
| port_speed 4Gb
| port_id 50050768014010FE
| port_status active
| port_speed 4Gb
| hardware 100
| iscsi_name iqn.1986-03.com.ibm:2145.tbcluster-12.node1
| iscsi_alias
| failover_active no
| failover_name node2
| failover_iscsi_name iqn.1986-03.com.ibm:2145.tbcluster-12.node2
| failover_iscsi_alias
| panel_name 01-1
| enclosure_id 1
| canister_id 1
| enclosure_serial_number 00001A2
| service_IP_address 192.168.70.121
| service_gateway 102.68.70.1
| service_subnet_mask 255.255.255.0
| service_IP_address_6
| service_gateway_6
| service_prefix_6

```

Isnodecandidate

The `Isnodecandidate` command lists all of the nodes that are available to add to the cluster.

Syntax

```

▶▶ svcinfo — — Isnodecandidate — — [ -nohdr ] [ -delim delimiter ] ▶▶

```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays a list of nodes that are available to add to the cluster. This includes nodes that are not already part of a cluster, but are compatible with the cluster software level. Nodes with hardware types that are incompatible with the installed software are not listed.

The following table describes the possible outputs:

Table 34. *lsnodecandidate* outputs

Attribute	Description
panel_name	SAN Volume Controller: Unique identifier for the nodes. Storwize V7000: A compound of enclosure serial number and canister ID, separated with a hyphen.
UPS_serial_number	SAN Volume Controller: The serial number of the UPS. Storwize V7000: Can be blank.
UPS_unique_id	SAN Volume Controller: The unique ID of UPS. Storwize V7000: Can be blank.
hardware	Describes the type of node.

Table 35. *lsnodecandidate* outputs

Attribute	Description
panel_name	A unique identifier for the nodes.
UPS_serial_number	The serial number of UPS.
UPS_unique_id	The unique ID of UPS.
hardware	Describes the type of node.

An invocation example

```
svcinfolsnodecandidate -delim :
```

The resulting output

```
id: panel_name: UPS_serial_number: UPS_unique_id: hardware
1: 146355: 10L3ASH: 202378I01C0D18D8: 8G4
```

Isnodedependentvdisks (Deprecated)

This command has been deprecated. Use the **Isdependentvdisks** command instead.

Isnodevpd

The **Isnodevpd** command displays the vital product data (VPD) for each node.

Syntax

```
svcinfolsnodevpd [-nohdr] [-delim delimiter] [node_id node_name]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each

column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. Using the **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

node_id | *node_name*

(Required) Specifies the node by its ID or name.

Description

This command displays the VPD for the specified node. Each field is reported on a new line. All fields are strings. The VPD is split into sections. Each section has a section heading. The number of fields in that section follows each heading. Each section is separated by an empty line.

For example:

```
section name:3 fields
field1:value
field2:value
field3:value
```

```
new section:x fields
```

```
...
```

Some sections contain information about multiple objects of that type. Each object within the section is separated by an empty line.

For example:

```
section name:4 fields
object1 field1:value
object1 field2:value
```

```
object2 field1:value
object2 field2:value
```

```
new section: x fields
```

```
...
```

Note: For 8F4, 8G4, and 8A4 nodes, the VPD displays the device serial number of the fibre-channel card as N/A.

An invocation example

```
svcinfo lsnodevpd 1
```

The resulting output

```
id 1
```

```
system board: 21 fields
part_number 43V7072
system_serial_number KD1438A
number_of_processors 4
number_of_memory_modules 6
number_of_fans 6
number_of_FC_cards 1
number_of_scsi/ide_devices 2
```

BIOS_manufacturer IBM Corp.
BIOS_version -[D6E124AUS-1.01]-
BIOS_release_date 04/30/2009
system_manufacturer IBM
system_product IBM System x -[2145CF8]-
version 00
planar_manufacturer IBM
planar_product 49Y6498
planar_version (none)
power_supply_part_number 39Y7201
CMOS_battery_part_number 33F8354
frame_assembly_part_number
ethernet_cable_part_number
service_processor_firmware 1.01

processor: 6 fields
processor_location Processor 1
manufacturer Intel(R) Corporation
version Intel(R) Xeon(R) CPU E5530 @ 2.40GHz
speed 2400
status Enabled
CPU_part_number 46D1266

memory module: 96 fields
part_number 44T1493
device_location DIMM01
bank_location BANK01
size (MB) No Module Installed
manufacturer Not Specified
serial_number Not Specified

part_number 44T1493
device_location DIMM02
bank_location BANK02
size (MB) 4096
manufacturer Samsung
serial_number 99062848

part_number 44T1493
device_location DIMM03
bank_location BANK03
size (MB) 4096
manufacturer Samsung
serial_number C7062848
...

fan: 12 fields
part_number 43V6929
location location1

part_number 43V6929
location location2

part_number 43V6929
location location3
...

Adapter card: 18 fields
card_type FC card
part_number 31P1337
port_numbers 1 2 3 4
location 0
device_serial_number 11S31P1333YM10MY96A206
manufacturer IBM
device QE8
card_revision 2
chip_revision 2.0


```
card_type SAS card
part_number 44E8690
port_numbers 1 2 3 4
location 0
device_serial_number 11S31P1299YM10MY948004
manufacturer IBMHUR
device Capri-PMC8001
card_revision Y
chip_revision 1.1
```

```
Fibre channel SFP: 48 fields
part_number 17P9211
manufacturer JDSU
device PLRXPLVCSH4921
serial_number C915EB06V
supported_speeds 2,4,8
connector_type LC
transmitter_type SN
wavelength 850
max_distance_by_cable_type OM1:20,OM2:50,OM3:150
hw_revision 1
port_number 1
WWPN 500507680140350d
...
```

```
device: 15 fields
part_number 31P1339
bus USB
device 0
model IBM USB Endeavour
revision 1.0
serial_number NA
approx_capacity 0
hw_revision 0
```

```
part_number 42D0673
bus scsi
device 0
model ST973452SS
revision B623
serial_number 3TA00BZ20109B623
approx_capacity 68
```

```
software: 8 fields
code_level 5.1.0.0 (build 16.1.0906240000)
node_name node1
ethernet_status 1
```

```
ethernet_status 0
WWNN 0x500507680100350d
id 1
MAC_address 00 21 5e 09 09 08
MAC_address 00 21 5e 09 09 0a
```

```
front panel assembly: 3 fields
part_number 31P1339
front_panel_id 161040
front_panel_locale en_US
```

```
UPS: 10 fields
electronics_assembly_part_number 64P8326
battery_part_number 31P0710
UPS_assembly_part_number 64P8326
input_power_cable_part_number CountryDependent
UPS_serial_number 1000840050
UPS_type 2I45UPS 1U
```


-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

ethernet_port_id

(Optional) Specifies the ID of an ethernet port, either 1 or 2. If omitted, a concise view is displayed for all ports. When you use this parameter, the detailed view of the specified port is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not use the *ethernet_port_id* parameter, the concise view displays all ports matching the filtering requirements that are specified by the **-filtervalue** parameter.

Description

This command lists all port IP addresses for each node in the cluster. The concise view displays two rows of output for each ethernet port. Each node has two ethernet ports.

Use the **svcinfolspportip** command with the optional **ethernet_port_id** parameter to display a detailed view of the specified port.

Both output rows for a port show the MAC address of that port if it can be determined. If the node and the ethernet link are online, the rows also show the speed and duplex state of the link. The duplex field can have values of Half or Full, or it is blank if the node is offline.

The first row for each port shows any iSCSI addresses that have been configured for that port and are not failed over to a different node. The failover field on this row is set to no. The second row for each port shows any iSCSI addresses that have been configured for the partner node, or for the local node with failover, and that are active on the port. The failover field on this row is set to yes.

The state field is set to unconfigured if there are no iSCSI addresses configured on the port. The state field is set to offline if there are configured addresses but the link is down, and online if the link is up. Any offline rows indicate a potential problem.

A concise invocation example

```
svcinfolspportip -delim :
```

The concise resulting output

```
id:node_id:node_name:IP_address:mask:gateway:IP_address_6:
prefix_6:gateway_6:MAC:duplex:state:speed:failover
1:1:dvt101794:9.71.47.129:255.255.254.0:9.71.46.1:::00:14:
  5e:33:51:92:Half:online:100Mb/s:no
1:1:dvt101794:::00:14:5e:33:51:92:Half:online:100Mb/s:yes
2:1:dvt101794:::00:14:5e:33:51:93::unconfigured::no
2:1:dvt101794:::00:14:5e:33:51:93::unconfigured::yes
1:2:dvt101760:9.71.47.83:255.255.254.0:9.71.46.1:::00:14:5e:
  7e:2a:58:Half:online:100Mb/s:no
1:2:dvt101760:::00:14:5e:7e:2a:58:Half:online:100Mb/s:yes
```

```

2:2:dvt101760:::::00:14:5e:7e:2a:59::unconfigured::no
2:2:dvt101760:::::00:14:5e:7e:2a:59::unconfigured::yes
1:3:dvt101761:9.71.47.253:255.255.254.0:9.71.46.1::::00:14:5e:
 33:50:fa:Half:online:100Mb/s:no
1:3:dvt101761:::::00:14:5e:33:50:fa:Half:online:100Mb/s:yes
2:3:dvt101761:::::00:14:5e:33:50:fb::unconfigured::no
2:3:dvt101761:::::00:14:5e:33:50:fb::unconfigured::yes
1:4:dvt101786:9.71.47.227:255.255.254.0:9.71.46.1::::00:14:5e:
 33:50:da:Half:online:100Mb/s:no
1:4:dvt101786:::::00:14:5e:33:50:da:Half:online:100Mb/s:yes
2:4:dvt101786:::::00:14:5e:33:50:db::unconfigured::no
2:4:dvt101786:::::00:14:5e:33:50:db::unconfigured::yes
1:5:destiny35:9.71.47.69:255.255.254.0:9.71.46.1::::00:21:5e:09:
 21:44:Full:online:1Gb/s:no
1:5:destiny35:::::00:21:5e:09:21:44:Full:online:1Gb/s:yes
2:5:destiny35:::::00:21:5e:09:21:46::unconfigured::no
2:5:destiny35:::::00:21:5e:09:21:46::unconfigured::yes
1:6:destiny34:9.71.46.239:255.255.254.0:9.71.46.1::::00:21:5e:09:
 21:54:Full:online:100Mb/s:no
1:6:destiny34:::::00:21:5e:09:21:54:Full:online:100Mb/s:yes
2:6:destiny34:::::00:21:5e:09:21:56::unconfigured::no
2:6:destiny34:::::00:21:5e:09:21:56::unconfigured::yes

```

A detailed invocation example

```
svcinfo lsportip 1
```

The detailed resulting output

```

id 1
node_id 1
node_name dvt101794
IP_address 9.8.7.1
mask 255.255.255.0
gateway 9.0.0.1
IP_address_6
prefix_6
gateway_6
MAC 00:14:5e:33:51:92
duplex Half
state online
speed 100Mb/s
failover no

```

lsquorum

| The **lsquorum** command lists the managed disks (MDisks) or drives that the cluster is currently using to
| store quorum data.

Syntax

```

▶▶ svcinfo — lsquorum — [ -nohdr ] [ -delim — delimiter ] [ quorum_index ] ▶▶

```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **delim** parameter overrides this behavior. Valid input for the **delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified character.

quorum_index

| (Optional) Specifies the quorum disk or drive by its index number. The number can be either 0, 1, or
| 2. When you use this parameter, a detailed view of the specified disk or drive is returned. If you do
| not specify a disk or drive, then a concise view of all quorum disks or drives is displayed.

Description

| This command displays a concise list or a detailed view of the MDisks or drives that the cluster is
| currently using to store quorum data. This information can be used to ensure that the quorum candidates
| are on separate storage subsystems.

| **Note:** The object type is either MDisk or drive. SAN Volume Controller uses only MDisks to hold
| quorum data. If the quorum object type is a drive, the controller ID and name fields are blank.

A concise invocation example

```
svcinfolsqorum
```

The concise resulting output

```
| quorum_index  status  id  name      controller_id  controller_name  active  object_type  
| 0             online  1  mdisk1   1             controller1      yes    mdisk  
| 1             online  2  mdisk2   1             controller1      no     mdisk  
| 2             online  33                no           drive
```

A detailed invocation example

```
svcinfolsqorum 1
```

The detailed resulting output


```
| quorum_index 2  
| status online  
| id 33  
| name  
| controller_id  
| controller_name  
| active no  
| object_type drive
```

lsrconsistgrp

The **lsrconsistgrp** command returns a concise list or a detailed view of Metro or Global Mirror consistency groups visible to the cluster.

Syntax

```
▶▶ svcinfo — — lsrconsistgrp — [ -filtervalue — attrib=value ] [ -nohdr ] →
```



Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- When using a wildcard, you must enclose the filter entry with double quotation marks (""), as follows:

```
svcinfolsrconsistgrp -filtervalue "name=md*"
```

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | object_name

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is displayed and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id | object_name* parameter, the concise view of all objects matching the filtering requirements that is specified by the **-filtervalue** parameter are displayed.

-filtervalue?

(Optional) Specifies that you want your report to display any or all of the list of valid filter attributes. The following filter attributes for the **svcinfolsrconsistgrp** command are valid:

- group_id
- name
- master_cluster_id
- master_cluster_name
- aux_cluster_id
- aux_cluster_name
- primary
- state
- relationship_count
- id

- copy_type

Description

This command returns a concise list or a detailed view of Metro or Global Mirror consistency groups that are visible to the cluster.

Table 36 provides possible values for the attributes that are displayed as data in the output views.

Table 36. `lsrcconsistgrp` command output values

Attribute	Value
primary	n/a, master, aux
state	inconsistent_stopped, inconsistent_copying, consistent_stopped, consistent_synchronized, idling, idling_disconnected, inconsistent_disconnected, consistent_disconnected, empty
freeze_time	The time in YY/MM/DD/HH/MM format.
status	online, primary_offline, secondary_offline
sync	in_sync, out_of_sync
copy_type	metro, global, empty_group

Note: The names of the Metro or Global Mirror relationships and consistency groups might be blank if the relationship or consistency groups are intercluster and the cluster partnership is disconnected.

- | The sync attribute has a value of in_sync when the contents are synchronized (identical) between VDisks (volumes). If write operations take place on either the primary or secondary volume after a consistent (stopped) or idling state occurs, they will no longer be synchronized.

A concise invocation example

```
svcinfolsrcconsistgrp -delim :
```

The concise resulting output

```
id:name:master_cluster_id:master_cluster_name:aux_cluster_id:aux_cluster_name:
primary:state:relationship_count:copy_type

248:jdemo_BA_cons1:0000020060406746:clusterB:0000020061413ABA:clusterA:master:
consistent_stopped:2:global
249:rccstgrp0:0000020061413ABA:clusterA:0000020061413ABA:clusterA::empty:0
:empty_group
250:jdemo_BA_cons2:0000020060406746:clusterB:0000020061413ABA:clusterA:master:
inconsistent_stopped:1:metro
251:BA_cons1:0000020060406746:clusterB:0000020061413ABA:clusterA:master:
consistent_stopped:4:metro
252:AB_cons2:0000020061413ABA:clusterA:0000020060406746:clusterB::empty:0
:empty_group
253:AB_cons1:0000020061413ABA:clusterA:0000020060406746:clusterB:aux:
consistent_stopped:3:global
254:AA_cons2:0000020061413ABA:clusterA:0000020061413ABA:clusterA::empty:0
:empty_group
255:AA_cons1:0000020061413ABA:clusterA:0000020061413ABA:clusterA:master:
consistent_synchronized:2:global
```

A detailed invocation example

```
svcinfolsrcconsistgrp -delim : 254
```

The detailed resulting output

```


id:254
name:rccstgrp0
master_cluster_id:0000010030A007E5
master_cluster_name:clusterA
aux_cluster_id:0000010030A007E5
aux_cluster_name:clusterA
primary:master
state:inconsistent_stopped
relationship_count:1
freeze_time:
status:online
sync:in_sync
copy_type:metro
RC_rel_id:2
RC_rel_name:aaa

```

lsrrelationship

The **lsrrelationship** command returns a concise list or a detailed view of Metro or Global Mirror relationships visible to the cluster.

Syntax



Parameters

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- When using a wildcard, you must enclose the filter entry with double quotation marks (" "), as follows:

```
svcinfo lsrrelationship -filtervalue "name=md*"
```

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed even if the **-nohdr** parameter is specified.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a

one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | *object_name*

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id* | *object_name* parameter, the concise view of all objects matching the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

-filtervalue?

(Optional) Specifies that you want your report to display any or all of the list of valid filter attributes. The valid filter attributes for the **svcinfo lsrelationship** command are:

- RC_rel_id
- RC_rel_name
- master_cluster_id
- master_cluster_name
- master_vdisk_id
- master_vdisk_name
- aux_cluster_id
- aux_cluster_name
- aux_vdisk_id
- aux_vdisk_name
- primary
- consistency_group_id
- consistency_group_name
- state
- progress
- copy_type

Description


This command returns a concise list or a detailed view of Metro or Global Mirror relationships visible to the cluster.

Table 37 provides possible values for the attributes that are displayed as data in the output views.

Table 37. lsrelationship command attributes and values

Attribute	Value
primary	n/a, master, aux
state	inconsistent_stopped, inconsistent_copying, consistent_stopped, consistent_synchronized, idling, idling_disconnected, inconsistent_disconnected, consistent_disconnected
progress	0-100, n/a
freeze time	The time in YY/MM/DD/HH/MM format.
status	online, primary_offline, secondary_offline
sync	n/a, in_sync, out_of_sync
copy_type	metro, global

Note: The names of the Metro or Global Mirror relationships and consistency groups can be blank if the relationship or consistency groups are inter-cluster and the cluster partnership is disconnected.



Parameters

-master *master_vdisk_id* | *master_vdisk_name*

(Required) Specifies a particular VDisk to use as the master VDisk. The command finds candidates that match the size of this VDisk. If you are requesting candidate VDIs on the local cluster, this command also matches the *io_group*.

-aux *aux_cluster_id* | *aux_cluster_name*

(Required) Specifies a remote cluster with VDisk candidates for an intercluster relationship. If you do not specify this parameter, the candidates on the local cluster are displayed.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays a list of VDIs that can be either the master or the auxiliary disk for a Metro or Global Mirror relationship. VDisk IDs and names are displayed.

An invocation example

```
svcinfolsrrelationshipcandidate -delim :
```


The resulting output

```
id:vdisk_name
0:vdisk0
4:vdisk4
```

lsrrelationshipprogress

You can use the **lsrrelationshipprogress** command to display the progress of the background copy of a Metro Mirror or Global Mirror relationship as a percentage. When the initial background copy process for a relationship has completed, null is displayed for the progress of that relationship.

Syntax





Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

rcrelationship_id | *rcrelationship_name*

Specifies the object ID or name of the specified type.

Description

This command displays the progress of the background copy of a Metro Mirror or Global Mirror relationship as a percentage.

An invocation example

```
svcinfolsrcrelationshipprogress -delim : 0
```



The resulting output

```
id:progress
0:58
```

Isrepairsevdiskcopyprogress

The `Isrepairsevdiskcopyprogress` command lists the repair progress for space-efficient VDisk copies.

Syntax



Parameters

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a

one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-copy *id*

(Optional) Lists the repair progress for the specified copy.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

vdisk_name | *vdisk_id*

(Optional) Specifies the virtual disk name or ID to list repair progress for. You must specify this parameter last on the command line. If you do not enter this parameter, the command lists progress for all space-efficient copies in the cluster.

Description

The **lsrepairvdiskcopyprogress** command lists the repair progress for space-efficient copies of the specified VDisk. If you do not specify a VDisk, the command lists the repair progress for all space-efficient copies in the cluster.

Note: Only run this command after running the **svctask repairsevdiskcopy** command, which you must only run as required by the Directed Maintenance Procedures or by IBM support.

An invocation example

```
svcinfo lsrepairsevdiskcopyprogress -delim :
```

The resulting output

```
| id:name:copy id:task:progress:estimated_completion_time
| 0:vdisk0:0:repairing:50:070301120000
| 0:vdisk0:1:repairing:51:070301120000
| 1:vdisk1:0:repairing:32:070301153500
```

An invocation example

```
svcinfo lsrepairsevdiskcopyprogress -delim : vdisk0
```

The resulting output

```
| id:name:copy id:task:progress:estimated_completion_time
| 0:vdisk0:0:repairing:50:070301120000
| 0:vdisk0:1:repairing:51:070301120000
```

An invocation example

```
svcinfo lsrepairsevdiskcopyprogress -delim : -copy 1 vdisk0
```


The resulting output

```
| id:name:copy id:task:progress:estimated_completion_time
| 0:vdisk0:1:repairing:51:070301120000
```

lsrepairvdiskcopyprogress

| The **lsrepairvdiskcopyprogress** command displays the progress of mirrored VDisk (volume) repairs and validations.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-copy *id*

(Optional) Specifies the VDisk copy ID to list repair progress for. If you do not specify this parameter, progress is displayed for all copies.

vdisk_name | *vdisk_id*

(Optional) Specifies the virtual disk name or ID to list repair progress for. You must specify this parameter last on the command line.

Description

The **lsrepairvdiskcopyprogress** command displays the progress of repairs and validations being made to mirrored VDIs. Use this command to track progress after running the **svctask repairvdiskcopy** command. You can specify a VDisk copy using the **-copy id** parameter. To display the VDIs that have two or more copies with an active task, specify the command with no parameters; it is not possible to have only one VDisk copy with an active task.

The command displays progress for the following types of VDisk copies:

- All VDisk copies display the same task; validate, medium or resync, depending on the specified parameter.
- All VDisk copies display the same percentage and estimated completion time.
- If specified, non-mirrored VDIs are displayed as a single copy with a blank task; they are not displayed in the full concise view.
- Once a task completes, the task is blank for all copies.
- If the task is blank, the percentage and the completion time are also blank.

An invocation example

```
svcinfo lsrepairvdiskcopyprogress -delim :
```

The resulting output

```
vdisk_id:vdisk_name:copy id:task:progress:estimated_completion_time
0:vdisk0:0:medium:50:070301120000
0:vdisk0:1:medium:50:070301120000
```

An invocation example

```
svcinfolrepairvdiskcopyprogress -delim : vdisk0
```

The resulting output

```
vdisk_id:vdisk_name:copy id:task:progress:estimated_completion_time
0:vdisk0:0:medium:50:070301120000
0:vdisk0:1:medium:50:070301120000
```

An invocation example

```
svcinfolsvdiskcopyrepairprogress -delim : -copy 0 vdisk0
```

The resulting output

```
vdisk_id:vdisk_name:copy id:task:progress:estimated_completion_time
0:vdisk0:0:medium:50:070301120000
```

lsvdiskdependentmaps

The **lsvdiskdependentmaps** command displays all FlashCopy mappings that must be stopped for the specified VDisk to be deleted.

Syntax

```
svcinfolsvdiskdependentmaps [-nohdr] [-delim delimiter] vdisk_name | vdisk_id
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

vdisk_name | vdisk_id

(Required) Specifies the name or ID of the VDisk for which the FlashCopy mappings are displayed.

Description

This command returns a list of the FlashCopy mappings that must be stopped before the specified VDisk can be deleted. Any mappings that are returned in the list for the VDisk are automatically stopped when the VDisk is deleted with the **force** option.

An invocation example

```
svcinfo lsrmvdiskdependentmaps -delim : 0
```

The resulting output

```
id:name  
2:fcmap2  
5:fcmap5
```

Isroute

The **Isroute** command displays the IP routing table.

Syntax

```
▶▶— svcinfo — — lsroute —————▶▶▶▶
```

Description

This command displays the IP routing table. The table provides details of the gateway that is used for IP traffic to a range of IP addresses for each ethernet port. This information can be used to diagnose configuration node accessibility problems. The **svcinfo Isroute** command is equivalent to the Linux **route** command.

An invocation example

```
svcinfo lsroute
```

The resulting output

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
9.71.46.0	0.0.0.0	255.255.254.0	U	0	0	0	eth0
127.0.0.0	0.0.0.0	255.0.0.0	U	0	0	0	lo
0.0.0.0	9.71.46.1	0.0.0.0	UG	0	0	0	eth0

Kernel IPv6 routing table

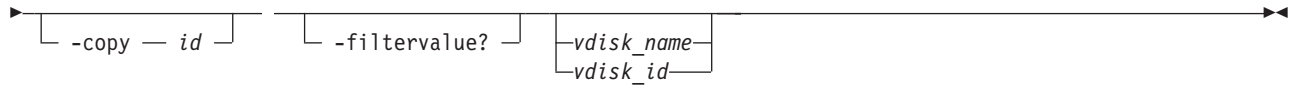
Destination	Next Hop	Flags	Metric	Ref	Use	Iface
2002:914:fc12:849::/64	::	UA	256	3675	0	eth0
fe80::/64	::	U	256	0	0	eth0
::/0	fe80::7:b4ff:fe00:500	UGDA	1024	1	0	eth0
::1/128	::	U	0	1441	1	lo
2002:914:fc12:849:214:5eff:fe33:5192/128	::	U	0	0	1	lo
fe80::214:5eff:fe33:5192/128	::	U	0	0	1	lo
ff00::/8	::	U	256	0	0	eth0

Issevdiskcopy

| The **Issevdiskcopy** command lists the space-efficient copies of the specified VDisks (volumes).

Syntax

```
▶▶— svcinfo — — lssevdiskcopy — [ -nohdr ] [ -bytes ] [ -delim — delimiter ] —▶▶▶▶
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-bytes

(Optional) Displays all capacities as bytes. Capacity values displayed in units other than bytes might be rounded.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-copy *id*

(Optional) Specifies the VDisk copy to list space-efficient copies for. You must specify a *vdisk_name* | *vdisk_id* value with this parameter.

-filtervalue?

(Optional) Displays a list of valid filter attributes. The following filters for the **svcinfolsevdiskcopy** command are valid:

- `mdisk_grp_id`
- `mdisk_grp_name`
- `overallocation`
- `autoexpand`
- `grainsize`

vdisk_name | *vdisk_id*

(Optional) Specifies the virtual disk name or ID to list space-efficient copies for. You must specify this parameter last on the command line. If you do not enter this parameter, the command lists all space-efficient VDisk (thin-provisioned volume) copies in the cluster.

Description

The **lssevdiskcopy** command lists all space-efficient copies of the specified VDisk. If you do not specify a VDisk, the command lists all space-efficient VDisk copies in the cluster.

The command provides a concise view of the space-efficient properties of the selected VDisk copies. Run the **svcinfolsvdiskcopy** command to see a concise view of the properties that are common to space-efficient and non-space-efficient VDisk copies. See the description of **svcinfolsvdisk** command for a description of the fields that is shown in the view.

The command returns values for the following VDisk copy attributes:

| **copy_id**
 | Specifies a system-assigned identifier for the VDisk copy. The value can be **0** or **1**.

| **status** The value can be **online** or **offline**. A copy is offline if all nodes cannot access the MDisk group that contains the copy.

| **sync** Indicates whether the VDisk copy is synchronized.

| **primary**
 | Indicates whether the VDisk copy is the primary copy. A VDisk has exactly one primary copy. The value can be **Yes** or **No**.

| **mdiskgrp_id/name**
 | Specifies the name and ID of the MDisk group that the VDisk copy belongs to.

| **type** Specifies the virtualization type of the VDisk. The value can be **striped**, **sequential** or **image**.

| **mdisk_id/name**
 | Specifies the MDisk that is used for sequential and image mode VDIsks.

| **fast_write_state**
 | Specifies the cache state of the VDisk copy. The value can be **empty**, **not_empty**, **corrupt**, or **repairing**. The value is always empty for non-space-efficient copies. A cache state of **corrupt** indicates that the VDisk is space-efficient and requires repair that is initiated by a **recovervdisk** command or the **repairsevdiskcopy** command.

| **used_capacity**
 | Specifies the portion of **real_capacity** that is being used to store data. For non-space-efficient copies, this value is the same as the VDisk capacity. If the VDisk copy is space-efficient, the value increases from zero to the **real_capacity** value as more of the VDisk is written to.

| **real_capacity**
 | Specifies the amount of physical storage that is allocated from an MDisk group to this VDisk copy. If the VDisk copy is not space-efficient, the value is the same as the VDisk capacity. If the VDisk copy is space-efficient, the value can be different.

| **free_capacity**
 | Specifies the difference between the **real_capacity** and **used_capacity** values.

| **overalllocation**
 | Expressed as a percentage, specifies the ratio of VDisk capacity to **real_capacity** values. This value is always **100** for non-space-efficient VDIsks.

| **autoexpand**
 | Specifies whether **autoexpand** is enabled on a space-efficient VDisk. The value can be **on** or **off**.

| **warning**
 | Expressed as a percentage, for space-efficient VDisk copies only. A warning is generated when the ratio of **used_capacity** to VDisk capacity reaches the specified level.

| **grainsize**
 | For space-efficient VDisk copies, specifies the grain size chosen for the VDisk copy when it was created.

| **se_copy**
 | Specifies if the copy is space-efficient.

| **easy_tier**
 | This value is set by the user and determines whether Easy Tier is permitted to manage the pool.

| **Note:**

- | 1. If **easy_tier** is *on*, then **easy_tier_status** can take on any value.
- | 2. if **easy_tier** is *off*, then **easy_tier_status** is *measured* or *inactive* .

| **easy_tier_status**
 | Which Easy Tier functions are active for the VDisk copy:

- | • *Active* : may move extents of this VDisk copy for performance (automatic data placement).
- | • *Measured*: statistics are being gathered for this VDisk copy, but no extents will be moved.
- | • *Inactive* : no Easy Tier function is active.

| **tier** Which tier information is being reported:

- | • generic_ssd
- | • generic_hdd

| **tier_capacity**
 | The total MDisk capacity assigned to the VDisk in the tier.

| **Note:** For space-efficient copies, the capacity by tier will be the real capacity.

| **An invocation example**

svcinfolssevdiskcopy -delim :

The resulting output

```
vdisk_id:vdisk_name:copy_id:mdisk_grp_id:mdisk_grp_name:capacity:used_capacity:real
_capacity:
free_capacity:overallocation:autoexpand:warning:grainsize
0:vv1:0:0:ppp:16.0GB:5.0MB:4.0GB:15.99GB:400:off:20:32
1:se1:0:0:ppp:16.0GB:1.0GB:4.0GB:15.00GB:400:off:20:32
1:se1:1:0:ppp:16.0GB:2.0GB:8.0GB:14.00GB:200:off:45:256
```

An invocation example

svcinfolssevdiskcopy -delim : -copy 0 0


The resulting output

```
| vdisk_id:0
| vdisk_name:vv1
| capacity:16.00GB
| copy_id:0
| status:online
| sync:yes
| primary:yes
| mdisk_grp:1
| mdisk_grp name:mdisk_group_1
| type:striped
| mdisk_id:
| mdisk_name:
| fast_write_state:not_empty
| used_capacity:2.00GB
| real_capacity:8.00GB
| free_capacity:6.00GB
| overallocation:200
| autoexpand:on
| warning:25
| grainsize:256
| se_copy:yes
| easy_tier:on
| easy_tier_status:active
| tier:generic_ssd
| tier_capacity:64.00MB
| tier:generic_hdd
| tier_capacity:7.94GB
```

lssnmppserver

The `lssnmppserver` command returns a concise list or a detailed view of SNMP servers that are configured on the cluster.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

snmp_server_name | *snmp_server_id*

(Optional) Specifies the name or ID of an existing SNMP server that must be listed.

Description

Use this command to display a concise list or a detailed view of SNMP servers that are configured on the cluster.

A concise invocation example

```
svcinfo lssnmppserver -delim :
```

The concise resulting output

```
id:name:IP_address:error:warning:info:port:community  
0:snmp0:192.135.60.4:on:on:on:78:public  
1:newserver:192.136.70.7:on:off:off:250:newcommunity
```

A detailed invocation example

```
svcinfo lssnmppserver snmp0
```

The detailed resulting output

```
id 0  
name snmp0  
IP_address 192.135.60.4  
error on
```

warning on
info on
port 78
community public

lssoftwaredumps (Deprecated)

| **Attention:** The `svcinfo lssoftwaredumps` command is deprecated. Use the `svcinfo lsdumps` command to
| display a list of files in a particular dumps directory.

lsoftwareupgradestatus

The `lsoftwareupgradestatus` command displays the status of a software upgrade.

Syntax

▶▶ `svcinfo` — — `lsoftwareupgradestatus` — — `--nohdr` ▶▶

Parameters

`--nohdr`
(Optional) Suppresses the display of headings.

Description

The `lsoftwareupgradestatus` command displays the status of a software upgrade.

Note: If a status of `stalled_non_redundant` is displayed, proceeding with the remaining set of node upgrades might result in offline VDIs. Contact an IBM service representative to complete the manual upgrade.

An invocation example

```
svcinfo lsoftwareupgradestatus
```

The resulting output

```
status  
upgrading
```

An invocation example

```
svcinfo lsoftwareupgradestatus
```


The resulting output

```
status  
stalled_non_redundant
```

lssyslogserver

The `lssyslogserver` command returns a concise list or a detailed view of syslog servers that are configured on the cluster.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

syslog_server_name | syslog_server_id

(Optional) Specifies the name or ID of an existing syslog server. When you use this parameter, a detailed view of the specified syslog server is returned. If you do not specify a syslog server name or ID, then a concise view of all syslog servers is displayed.

Description

Use this command to display a concise list or a detailed view of syslog servers that are configured on the cluster.

A concise invocation example

```
svcinfo lssyslogserver -delim :
```

The concise resulting output

```
id:name:IP_address:facility:error:warning:info
0:syslog0:192.135.60.4:0:on:on:on
1:newserver:192.136.70.7:4:on:off:off
```

A detailed invocation example

```
svcinfo lssyslogserver 0
```

The detailed resulting output

```
id 0
name syslog0
IP_address 192.135.60.4
facility 0
error on
warning on
info on
```

lstimezones

The **lstimezones** command lists the time zones that are available on the cluster. Each timezone is assigned an ID that can be used in the **svctask settimezone** command to set the time zone.

Syntax

```
➤— svcinfo — — lstimezones — — [ -nohdr ] [ -delim — delimiter ] —————➤
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by a colon character.

Description

This command displays a list of all the time zones that are available on the cluster. Each time zone is assigned an ID. This ID can be used in the **svctask settimezone** command.

An invocation example

```
svcinfo lstimezones
```

The resulting output

```
id timezone
0 Africa/Abidjan
1 Africa/Accra
2 Africa/Addis_Ababa
3 Africa/Algiers
4 Africa/Asmera
5 Africa/Bamako
6 Africa/Bangui
```

lsuser

Use the **lsuser** command to display a list of the users that have been created on the cluster.

Syntax

```
➤— svcinfo — — lsuser — — [ -nohdr ] [ -delim — delimiter ] —————➤
```

`-filtervalue — attribute=value`

`-filtervalue?`

`userid_or_name`

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is the asterisk (`*`), which must be used as the first or last character in the string.
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (`"`), as follows:
`svcinfo lsuser -filtervalue "usergrp_name=md*"`

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** *attribute=value* parameter:

- password
- ssh_key
- remote
- usergrp_id
- usergrp_name

userid_or_name

(Optional) Specifies the ID or name of the user for which the association is being deleted. If this is specified, the detailed view for the specified user is displayed in the output. If you do not specify an ID or name, the concise view is displayed.

Description

This command displays a list of users that have been created on the cluster.

An invocation example

```
svcinfo lsuser
```


The resulting output

id	name	password	ssh_key	remote	usergrp_id	usergrp_name
0	superuser	yes	no	no	0	SecurityAdmin
1	simon	no	yes	no	2	CopyOperator
2	jane	yes	no	no	3	Service
3	kip	yes	yes	yes		

lsusergrp

Use the **lsusergrp** command to display a list of the user groups that have been created on the cluster.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- The wildcard character is the asterisk (*), which must be used as the first or last character in the string.
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:

```
svcinfo lsusergrp -filtervalue "role=md*"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue *attribute=value*** parameter:

- role
- remote

usergrp_id_or_name

(Optional) Specifies the ID or name of the user group to view. If you do not specify an ID or name, all groups are displayed.

Description

This command displays a list of user groups that have been created on the cluster.

An invocation example

```
svcinfolusergrp
```

The resulting output

id	name	role	remote
0	SecurityAdmin	SecurityAdmin	yes
1	Administrator	Administrator	no
2	CopyOperator	CopyOperator	no
3	Service	Service	yes
4	Monitor	Monitor	no
5	support	Service	no

lsvdisk

The **lsvdisk** command displays a concise list or a detailed view of VDisks that are recognized by the cluster.

Syntax

```
svcinfolsvdisk [-filtervalue attrib=value] [-nohdr] [-bytes]
                 [-delim delimiter] [-filtervalue?] [object_id object_name]
```

Parameters

-filtervalue *attrib=value*

(Optional) Specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are displayed. If a capacity is specified, the units must also be included.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-bytes

(Optional) Displays all capacities as bytes. Capacity values displayed in units other than bytes might be rounded. When filtering on capacity, use a unit of bytes, **-unit b**, for exact filtering. For space efficient copies, the capacity by tier will be the real capacities.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

object_id | *object_name*

(Optional) Specifies the name or ID of an object. When you use this parameter, the detailed view of the specific object is returned and any value that is specified by the **-filtervalue** parameter is ignored. If you do not specify the *object_id* | *object_name* parameter, the concise view of all objects matching the filtering requirements that is specified by the **-filtervalue** parameter are displayed.

-filtervalue?

(Optional) Displays a list of valid filter attributes. The following filters for the **svcinfo lsvdisk** command are valid:

- vdisk_name
- vdisk_id
- vdisk_UID
- fc_map_count
- copy_count
- IO_group_id
- IO_group_name
- status
- mdisk_grp_name
- mdisk_grp_id
- capacity
- type
- FC_id
- FC_name
- RC_id
- RC_name
- name
- id

Note: It is not possible to filter the **svcinfo lsvdisk** command with `mdisk_grp_name=many` to identify mirrored VDIs. Instead, filter on `copy_count=2`.

Description

This command displays a concise list or a detailed view of attributes for all VDIs and VDisk copies in the cluster.

The VDisk is offline and unavailable if one of the following takes place:

- Both nodes in the I/O group are missing.
- None of the nodes in the I/O group that are present can access the VDisk.
- All synchronized copies for this VDisk are in MDisk groups that are offline.
- The VDisk is formatting.

If you have a degraded VDisk and all of the associated nodes and MDIs are online, call the IBM Support Center for assistance. A VDisk is reported as degraded if any of the following occurs:

- One of the nodes in the I/O group is missing.
- One of the nodes in the I/O group cannot access all the MDIs in the MDisk group that the VDisk spans. In this case MDIs are shown as degraded and the DMPs for MDIs should be followed to resolve the problem.
- The fast write cache pins data for one or more VDIs in the I/O group and is unable to perform a failback until the situation is resolved. An error log indicating that the cache has pinned data is

displayed. Follow the directed maintenance procedures for this error log to resolve the problem. The most common causes of pinned data are the following:

- One or more VDisks in an I/O group is offline due to an asymmetric failure and has pinned data in the cache. Asymmetric failures can occur because of SAN fabric faults or misconfiguration, back-end controller faults or misconfiguration or because repeated errors has lead to the cluster excluding access to a MDisk through one or more nodes.
- One or more VDisks in an I/O group is offline due to a problem with a FlashCopy mapping.

The command returns values for the following VDisk attributes:

IO_groups_id/name

Specifies the I/O Group that the VDisk belongs to.

status The value can be **online**, **offline** or **degraded**.

mdisk_grp_id/name

Specifies the name and ID of the MDisk group that the VDisk belongs to. If the VDisk has more than one copy, these fields display **many**.

type Specifies the virtualization type of the VDisk. The value can be **striped**, **sequential**, **image** or **many**. The value **many** indicates that the VDisk has more than one copy, which can have different virtualization types.

capacity

Specifies the total capacity of the VDisk.

formatted

Indicates whether the VDisk was formatted when it was created. The value can be **Yes** or **No**.

mdisk_id/name

Specifies the MDisk that is used for sequential and image mode VDisks. If the VDisk has more than one copy, these fields display **many**.

FC_id/name

Specifies the name and ID of the FlashCopy mapping that the VDisk belongs to. The value **many** indicates that the VDisk belongs to more than one FlashCopy mapping.

RC_id/name

Specifies the name and ID of the Global Mirror or Metro Mirror relationship that the VDisk belongs to.

vdisk_UID

Specifies the UID of the VDisk.

throttling

Specifies the throttle rate of the VDisk.

preferred_node_id

Specifies the ID of the preferred node for the VDisk.

fast_write_state

Specifies the cache state for the VDisk. The value can be **empty**, **not_empty**, **corrupt**, or **repairing**. A cache state of **corrupt** indicates that the VDisk requires recovery by using one of the **recovervdisk** commands. A cache state of **repairing** indicates that repairs initiated by a **recovervdisk** command are in progress.

cache Specifies the cache mode of the VDisk. The value can be **readwrite** or **none**.

udid Specifies the unit number for the VDisk. Only OpenVMS hosts require a unit number.

fc_map_count

Specifies the number of FlashCopy mappings that the VDisk belongs to.

sync_rate
Specifies the rate for synchronization for mirrored copies.

| **se_copy_count**
| Indicates if the copy is space-efficient:
| • yes
| • no

The command returns values for the following VDisk copy attributes:

copy_id
Specifies a system-assigned identifier for the VDisk copy. The value can be **0** or **1**.

status The value can be **online** or **offline**. A copy is offline if all nodes cannot access the MDisk group that contains the copy.

sync Indicates whether the VDisk copy is synchronized.

primary
Indicates whether the VDisk copy is the primary copy. A VDisk has exactly one primary copy. The value can be **Yes** or **No**.

mdiskgrp_id/name
Specifies the name and ID of the MDisk group that the VDisk copy belongs to.

type Specifies the virtualization type of the VDisk. The value can be **striped**, **sequential** or **image**.

mdisk_id/name
Specifies the MDisk that is used for sequential and image mode VDIsks.

fast_write_state
Specifies the cache state of the VDisk copy. The value can be **empty**, **not_empty**, **corrupt**, or **repairing**. The value is always empty for non-space-efficient copies. A cache state of **corrupt** indicates that the VDisk is space-efficient and requires repair that is initiated by a **recovervdisk** command or the **repairsevdiskcopy** command.

used_capacity
Specifies the portion of **real_capacity** that is being used to store data. For non-space-efficient copies, this value is the same as the VDisk capacity. If the VDisk copy is space-efficient, the value increases from zero to the **real_capacity** value as more of the VDisk is written to.

real_capacity
Specifies the amount of physical storage that is allocated from an MDisk group to this VDisk copy. If the VDisk copy is not space-efficient, the value is the same as the VDisk capacity. If the VDisk copy is space-efficient, the value can be different.

free_capacity
Specifies the difference between the **real_capacity** and **used_capacity** values.

overalllocation
Expressed as a percentage, specifies the ratio of VDisk capacity to **real_capacity** values. This value is always **100** for non-space-efficient VDIsks.

autoexpand
Specifies whether **autoexpand** is enabled on a space-efficient VDisk. The value can be **on** or **off**.

warning
Expressed as a percentage, for space-efficient VDisk copies only. A warning is generated when the ratio of **used_capacity** to VDisk capacity reaches the specified level.

grainsize
For space-efficient VDisk copies, specifies the grain size chosen for the VDisk copy when it was created.

| **se_copy**
 | Specifies if the copy is space-efficient.

| **easy_tier**
 | This value is set by the user and determines whether Easy Tier is permitted to manage the pool.

| **Note:**
 | 1. If *easy_tier* is *on*, then *easy_tier_status* can take on any value.
 | 2. if *easy_tier* is *off*, then *easy_tier_status* is *measured* or *inactive* .

| **easy_tier_status**
 | Which Easy Tier functions are active for the VDisk copy:
 | • *Active* : may move extents of this VDisk copy for performance (automatic data placement).
 | • *Measured*: statistics are being gathered for this VDisk copy, but no extents will be moved.
 | • *Inactive* : no Easy Tier function is active.

| **tier** The tier information being reported:
 | • *generic_ssd*
 | • *generic_hdd*

| **tier_capacity**
 | The total MDisk capacity assigned to the VDisk in the tier.

| **Note:** For space-efficient copies, the capacity by tier will be the real capacity.

| **A concise invocation example**

| `svcinfolsvdisk -delim :`

| **The resulting output**

```
| id:name:IO_group_id:IO_group_name:status:mdisk_grp_id:mdisk_grp_name:capacity:type:FC_id:
| FC_name:RC_id:RC_name:vdisk_UID:fc_map_count:copy_count:fast_write_state:se_copy_count
| 0:vdisk0:0:io_grp0:degraded:0:mdiskgrp0:10.00GB:striped:::::60050768018300003000000000000000:0:1
| :empty:0
```

| **A detailed invocation example for a VDisk**

```
| svcinfolsvdisk -delim : vv1
| id:0
| name:vv1
| IO_group_id:0
| IO_group_name:io_grp0
| status:degraded
| mdisk_grp_id:many
| mdisk_grp_name:many
| capacity:16.00GB
| type:many
| formatted:no
| mdisk_id:many
| mdisk_name:many
| FC_id:
| FC_name:
| RC_id:
| RC_name:
| vdisk_UID:0000000000000000AB:6005076801CF003F2800000000000000
| throttling:0
| preferred_node_id:1
| fast_write_state:empty
| cache:readwrite
| udid:1234
```

```

| fcmap_count:0
| sync_rate:25
| copy_count:2
| se_copy_count:0
|
| copy_id:0
| status:online
| sync:yes
| primary:yes
| mdisk_grp:1
| mdisk_grp_name:mdisk_group_1
| type:striped
| mdisk_id:
| mdisk_name:
| fast_write_state:empty
| used_capacity:2.00GB
| real_capacity:8.00GB
| free_capacity:6.00GB
| overallocation:200
| autoexpand:on
| warning:25
| grainsize:256
| se_copy:yes
| easy_tier:off
| easy_tier_status:inactive
| tier:generic_ssd
| tier_capacity:0.00MB
| tier:generic_hdd
| tier_capacity:8.00GB
|
| copy_id:1
| status:offline
| sync:no
| primary:no
| mdisk_grp:2
| mdisk_grp_name:mdisk_group_2
| type:striped
| mdisk_id:
| mdisk_name:
| fast_write_state:not_empty
| used_capacity:2.00GB
| real_capacity:4.00GB
| free_capacity:2.00GB
| overallocation:400
| autoexpand:on
| warning:20
| grainsize:256
| se_copy:yes
| easy_tier:on
| easy_tier_status:active
| tier:generic_ssd
| tier_capacity:64.00MB
| tier:generic_hdd
| tier_capacity:3.94GB

```

lsvdiskcopy


| The `lsvdiskcopy` command lists VDisk (volume) copy information.

Syntax

```

| ►► svcinfo — lsvdiskcopy — [ -nohdr ] [ -bytes ] [ -delim delimiter ] —►
|

```



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-bytes

(Optional) Displays all capacities as bytes. Capacity values displayed in units other than bytes might be rounded.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-copy copy_id

(Optional) Specifies the VDisk copy to list information for. You must specify a *vdisk_name* | *vdisk_id* value with this parameter.

-filtervalue?

(Optional) Displays a list of valid filter attributes. The following filters for the **svcinfolsvdiskcopy** command are valid:

- primary
- status
- sync
- mdisk_grp_id
- mdisk_grp_name
- type
- easy_tier
- easy_tier_status

vdisk_name | *vdisk_id*

(Optional) Specifies the VDisk to list copy information for. You must specify this parameter last on the command line. If you specify a *vdisk_name* | *vdisk_id* value only, all copies for the VDisk are listed.

Description

The **lsvdiskcopy** command lists information for VDisk copies. If you specify the command with no parameters, all VDIsks and copies in the cluster are listed.

The command returns values for the following VDisk copy attributes:

copy_id

Specifies a system-assigned identifier for the VDisk copy. The value can be **0** or **1**.

status The value can be **online** or **offline**. A copy is offline if all nodes cannot access the MDisk group that contains the copy.

sync Indicates whether the VDisk copy is synchronized.

primary

Indicates whether the VDisk copy is the primary copy. A VDisk has exactly one primary copy. The value can be **Yes** or **No**.

mdiskgrp_id/name

Specifies the name and ID of the MDisk group that the VDisk copy belongs to.

type Specifies the virtualization type of the VDisk. The value can be **striped**, **sequential** or **image**.

mdisk_id/name

Specifies the MDisk that is used for sequential and image mode VDIs.

fast_write_state

Specifies the cache state of the VDisk copy. The value can be **empty**, **not_empty**, **corrupt**, or **repairing**. The value is always empty for non-space-efficient copies. A cache state of **corrupt** indicates that the VDisk is space-efficient and requires repair that is initiated by a **recovervdisk** command or the **repairsevdiskcopy** command.

used_capacity

Specifies the portion of **real_capacity** that is being used to store data. For non-space-efficient copies, this value is the same as the VDisk capacity. If the VDisk copy is space-efficient, the value increases from zero to the **real_capacity** value as more of the VDisk is written to.

real_capacity

Specifies the amount of physical storage that is allocated from an MDisk group to this VDisk copy. If the VDisk copy is not space-efficient, the value is the same as the VDisk capacity. If the VDisk copy is space-efficient, the value can be different.

free_capacity

Specifies the difference between the **real_capacity** and **used_capacity** values.

overalllocation

Expressed as a percentage, specifies the ratio of VDisk capacity to **real_capacity** values. This value is always **100** for non-space-efficient VDIs.

autoexpand

Specifies whether **autoexpand** is enabled on a space-efficient VDisk. The value can be **on** or **off**.

warning

Expressed as a percentage, for space-efficient VDisk copies only. A warning is generated when the ratio of **used_capacity** to VDisk capacity reaches the specified level.

grainsize

For space-efficient VDisk copies, specifies the grain size chosen for the VDisk copy when it was created.

| **se_copy**

| Specifies if the copy is space-efficient.

| **easy_tier**

| This value is set by the user and determines whether Easy Tier is permitted to manage the pool.

| **Note:**

| 1. If **easy_tier** is *on*, then **easy_tier_status** can take on any value.

| 2. If **easy_tier** is *off*, then **easy_tier_status** is *measured* or *inactive*.

| **easy_tier_status**

| Which Easy Tier functions are active for the VDisk copy:

- | *Active* : may move extents of this VDisk copy for performance (automatic data placement).
- | *Measured*: statistics are being gathered for this VDisk copy, but no extents will be moved.
- | *Inactive* : no Easy Tier function is active.

| **tier** Which tier information is being reported:

- | `generic_ssd`
- | `generic_hdd`

| **tier_capacity**

| The total MDisk capacity assigned to the VDisk in the tier.

| **Note:** For space-efficient copies, the capacity by tier will be the real capacity.

An invocation example

```
svcinfolsvdiskcopy -delim :
```

The resulting output

```
vdisk_id:vdisk_name:copy_id:status:sync:primary:mdisk_grp_id:mdisk_grp_name:
capacity:type:se_copy:easy_tier:easy_tier_status
0:RAM_V2:0:online:yes:yes:2:RAM_MDG2:5.00GB:striped:yes:on:inactive
1:RAM_V3:0:online:yes:yes:2:RAM_MDG2:5.00GB:striped:no:on:inactive
2:RAM_V4:0:online:yes:yes:1:RAM_MDG3:5.00GB:striped:no:on:inactive
3:RAM_V5:0:online:yes:yes:2:RAM_MDG2:5.00GB:striped:yes:on:inactive
3:RAM_V5:1:online:yes:no:2:RAM_MDG2:5.00GB:striped:yes:on:inactive
4:RAM_V1:0:online:yes:yes:3:RAM_MDG1:5.00GB:striped:no:on:inactive
5:RAM_V6:0:online:yes:yes:0:RAM_MDG4:5.00GB:striped:yes:on:inactive
```

An invocation example

```
svcinfolsvdiskcopy -copy 0 -delim : vv1
```

The resulting output

```
vdisk_id:0
vdisk_name:vv1
capacity:16.00GB
copy_id:0
status:online
sync:yes
primary:yes
mdisk_grp:1
mdisk_grp_name:mdisk_group_1
type:striped
mdisk_id:
mdisk_name:
fast_write_state:not_empty
used_capacity:2.00GB
real_capacity:8.00GB
free_capacity:6.00GB
overallocation:200
autoexpand:on
warning:25
grain_size:256
se_copy:yes
easy_tier:on
easy_tier_status:active
tier:generic_ssd
tier_capacity:64.00MB
tier:generic_hdd
tier_capacity:7.94GB
```

lsvdiskdependentmaps

The `lsvdiskdependentmaps` command displays all FlashCopy mappings with target virtual disks (VDisks) that are dependent upon data held on the specified VDisk.

Syntax

```
svcinfolsvdiskdependentmaps [vdisk_id | vdisk_name]
```

Parameters

vdisk_id | *vdisk_name*

(Required) Specifies the name or ID of a virtual disk (VDisk).

Description

The `lsvdiskdependentmaps` command displays FlashCopy mappings that have target VDisks that are dependent upon data held on the specified *vdisk_id* | *vdisk_name*. This can be used to determine whether a FlashCopy mapping can be prepared. Issue the command for the target VDisk *vdisk_id* | *vdisk_name* of the FlashCopy mapping to be prepared. If no FlashCopy mappings are returned, the FlashCopy mapping can be prepared. Any FlashCopy mappings that are returned in the list must be stopped or be in the `idle_or_copied` state, before the new FlashCopy mapping can be prepared.

A concise invocation example

```
svcinfolsvdiskdependentmaps -delim : 0
```

The concise resulting output

```
id:name
2:fcmap2
5:fcmap5
```

lsvdiskextent

The `lsvdiskextent` command lists the MDisk extents that are provided for the specified VDisks.

Syntax

```
svcinfolsvdiskextent [-copy copy_id] [-nohdr]
                        [-delim delimiter] [vdisk_name | vdisk_id]
```

Parameters

`-copy` *copy_id*

(Optional) Displays a list of MDisks that are members of the specified VDisk copy.

`-nohdr`

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The `-nohdr` parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

vdisk_name | *vdisk_id*

(Required) Specifies one or more virtual disk IDs or names.

Description

The **lsvdiskextent** command displays a list of MDisk IDs and the number of extents that each MDisk provides to the specified VDisks.

Each VDisk is constructed from one or more MDisks. To determine the relationship between a VDisk and its MDisks, issue the following command:

```
svcinfolsvdiskmember vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the VDisk. This command displays a list of MDisk IDs that make up the VDisk.

To determine the number of extents that are provided by each MDisk, issue the following command:

```
svcinfolsvdiskextent vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the VDisk. This command displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the given VDisk.

To determine the relationship between MDisks and VDisks, issue the following command for each MDisk:

```
svcinfolsmdiskmember mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This command displays a list of IDs that corresponds to the VDisks that are using this MDisk.

To determine the relationship between MDisks and VDisks, and the number of extents that are used by each VDisk, you must use the command-line interface. For each MDisk, issue the following command:

```
svcinfolsmdiskextent mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This command displays a table of VDisk IDs and the corresponding number of extents that are used by each VDisk.

An invocation example

```
svcinfolsvdiskextent -delim : vdisk0
```

The resulting output

```
id:number_extents  
0:0
```

lsvdiskfcmappcopies

The `lsvdiskfcmappcopies` command displays a list of all FlashCopy mappings with a target VDisk containing a valid copy of the specified VDisk.

Syntax

```
svcinfo -- lsvdiskfcmappcopies [-nohdr] [-delim delimiter] vdisk_name | vdisk_id
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (`:`) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

vdisk_name | vdisk_id

(Required) Specifies the name or ID of the VDisk for which the FlashCopy mappings are displayed.

Description

This command returns a list of the FlashCopy mappings that have a target VDisk with a valid copy of the specified VDisk. The target VDIsks of these mappings can be considered as candidate source VDIsks for mappings to restore from.

The mappings returned are in the copying, idle_copied, or stopping state with 100% progress.

An invocation example

```
svcinfo lsvdiskfcmappcopies -delim : 0
```

The resulting output

```
id:name:status:progress:difference:start_time:target_vdisk_id:
target_vdisk_name:group_id:group_name
2:fcmap2:copying:80:10:060627083137:10:vdisk10::
5:fcmap5:idle_copied:100:20:060627073130:12:vdisk12:1:fccstgrp1
```

lsvdiskfcmappings

The `lsvdiskfcmappings` command displays a list of FlashCopy mappings to which the VDisk belongs. A VDisk can be part of up to 256 FlashCopy mappings.

Syntax

```
▶▶ svcinfo — — lsvdiskfcmappings — [ vdisk_name ] —————▶▶
                                   [ vdisk_id ]
```

Parameters

vdisk_name | *vdisk_id*

(Required) Specifies the name or ID of the VDisk for which a list of all FlashCopy mappings is required.

Description

The **lsvdiskfcmappings** command returns a list of all FlashCopy mappings that the VDisk is a member of. The list is returned in no particular order.

An invocation example

```
svcinfo lsvdiskfcmappings -delim : vdisk2
```

The resulting output

```
fc_id:fc_name
1:fcmap1
3:fcmap3
```

lsvdiskhostmap

Use the **lsvdiskhostmap** command to list the VDisk to the host mapping. These hosts have the specified virtual disk mapped to them; the virtual disk is visible to these hosts.

Syntax

```
▶▶ svcinfo — — lsvdiskhostmap — [ -nohdr ] [ -delim — delimiter ] —————▶▶
▶ [ vdisk_id ] —————▶▶
  [ vdisk_name ]
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

vdisk_id | *vdisk_name*

(Required) Specifies the ID or name of the virtual disk. The cluster displays a list of all the hosts to which this virtual disk is mapped and the SCSI ID by which the virtual disk is mapped.

Description

This command displays a list of host IDs and names. These hosts have the specified virtual disk mapped to them; that is, the virtual disk is visible to these hosts. The SCSI LUN ID is also displayed. The SCSI LUN ID is the ID by which the virtual disk is recognized by the host.

Determining the host that a VDisk is mapped to: List the hosts that this VDisk is mapped to, by issuing the following command:

```
svcinfolsvdiskhostmap vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the VDisk. A list is displayed. Look for the host name or ID to determine which host this VDisk is mapped to. If no data is displayed, the VDisk is not mapped to any hosts.

An invocation example

```
svcinfolsvdiskhostmap bbb
```

The resulting output

id	name	SCSI_id	host_id	host_name	vdisk_UID
200	bbb	0	9	mchost13	600507680197014B000000000000002A0

lsvdisklba

The **lsvdisklba** command lists the VDisk (volume) and logical block address (LBA) for the specified MDisk LBA.

Syntax

```
svcinfolsvdisklba -mdisklba mdisklba [-delim delimiter]
```

```
svcinfolsvdisklba [-nohdr] -mdisk [mdisk_id | mdisk_name]
```

Parameters

-mdisklba *mdisklba*

(Required) Specifies the 64-bit hexadecimal LBA on the MDisk.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a

one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

-mdisk *mdisk_id* | *mdisk_name*
 (Required) Specifies the MDisk name or ID.

Description

The `lsvdisklba` command returns the LBA of the VDisk that is associated with the MDisk LBA.

If applicable, the command also lists the range of LBAs on both the VDisk and MDisk that are mapped in the same extent, or for space-efficient disks, in the same grain.

Table 38 provides command output that depends on several variables.

Table 38. *lsvdisklba* command output scenarios

Field	Typical scenario	Quorum disk	Space-efficient metadata	Extent not allocated	Formatting extent	Extent allocated to space-efficient disk, LBA not used on space-efficient disk
<code>copy_id</code>	yes	no	yes	no	yes	yes
<code>vdisk_id</code>	yes	no	yes	no	yes	yes
<code>vdisk_name</code>	yes	no	yes	no	yes	yes
<code>type</code>	allocated	metadata	metadata	unallocated	formatting	unallocated
<code>vdisk_lba</code>	yes	no	no	no	no	no
<code>vdisk_start</code>	yes	no	no	no	no	no
<code>vdisk_end</code>	yes	no	no	no	no	no
<code>mdisk_start</code>	yes	yes	yes	yes	yes	yes
<code>mdisk_end</code>	yes	yes	yes	yes	yes	yes

An invocation example

```
svcinfolsvdisklba -mdisk 1 -mdisklba 0x0
```


The resulting output

```
vdisk_id
vdisk_name
copy_id
type metadata
vdisk_lba 0x00090000
vdisk_start
vdisk_end
mdisk_start 0x00000000
| mdisk_end 0x1FFFFFFFFF
```

lsvdiskmember

The `lsvdiskmember` command displays a list of MDisks that are members of the specified VDisk.

Syntax



Parameters

-copy *copy_id*

(Optional) Displays a list of MDisks that are members of the specified VDisk copy.

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter.

vdisk_id | vdisk_name

(Required) Displays a list of MDisks that are members of the specified VDisk.

Description

This command displays a list of managed disks, which provide extents that make up the virtual disk that is specified by the ID.

Every VDisk is constructed from one or more MDisks. At times, you might have to determine the relationship between the two objects. The following procedure allows you to determine the relationships.

If you use the **svcinfo lsvdiskmember** command, the concise view displays a list of virtual disks. These are the virtual disks that are using extents on the managed disk that is specified by the ID. The list displays the members of the respective object and is independent of the state of the individual members; that is, if they are in offline state, they are still displayed.

To determine the relationship between VDIsks and MDisks, issue the following command:

```
svcinfo lsvdiskmember vdisk_id | vdisk_name
```

where *vdisk_id | vdisk_name* is the name or ID of the VDisk. This displays a list of IDs that correspond to the MDisks that make up the VDisk.

To determine the relationship between VDIsks and MDisks, and the number of extents that are provided by each MDisk, you must use the command-line interface. Issue the following command:

```
svcinfo lsvdiskextent vdisk_id | vdisk_name
```

where *vdisk_id | vdisk_name* is the name or ID of the VDisk. This displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the specified VDisk.

To determine the relationship between MDisks and VDIs, issue the following command:

```
svcinfolsmdiskmember mdisk_id | mdisk_name
```

where *mdisk_id* | *mdisk_name* is the name or ID of the MDisk. This displays a list of IDs that correspond to the VDIs that are using this MDisk.

To determine the relationship between MDisks and VDIs, and the number of extents that are used by each VDisk, you must use the command-line interface. For a specified MDisk, issue the following command:

```
svcinfolsmdiskextent mdisk_id | mdisk_name
```

where *mdisk_id* | *mdisk_name* is the name or ID of the MDisk. This displays a table of VDisk IDs and the corresponding number of extents that are used by each VDisk.

An invocation example

```
svcinfolsvdiskmember 1
```

The resulting output

```
id  
2
```

lsvdiskprogress

The `lsvdiskprogress` command tracks the progress during new virtual disk formatting.

Syntax

```
svcinfolsvdiskprogress [-nohdr] [-delim delimiter] vdisk_id | vdisk_name
```

Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim *delimiter*

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter `-delim :` on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by a colon character.

vdisk_id | *vdisk_name*

(Optional) Specifies the VDisk ID or name. If you do not specify this parameter, the progress of all VDIs currently being formatted is displayed.

Description

This command displays the progress of the format of a new virtual disk as a completed percentage. If the VDisk has multiple copies, the command reports the average progress of the format.

An invocation example

```
svcinfo lsvdiskprogress -delim : 0
```

The resulting output

```
id:progress  
0:58
```

lsvdisksyncprogress

The `lsvdisksyncprogress` command displays the progress of VDisk copy synchronization.

Syntax

```
▶▶—svcinfo— —lsvdisksyncprogress— — -copy — id — — vdisk_name —  
— vdisk_id —▶▶
```

Parameters

-copy *id*

(Optional) Specifies the VDisk copy ID to list synchronization progress for. You must also specify a *vdisk_name* | *vdisk_id* value. If you do not specify this parameter, progress is displayed for all copies.

vdisk_name | *vdisk_id*

(Optional) Specifies the virtual disk name or ID to list synchronization progress for.

Description

The `lsvdisksyncprogress` command displays the progress of VDisk copy synchronization. To display the VDisk copies that require synchronization, specify the command with no parameters. Estimated completion time is in the YYMMDDHHMMSS format. The command displays progress for the following types of VDIsks:

- A synchronized copy displays a progress of 100 and a blank estimated completion time.
- An offline copy displays a blank estimated completion time, and a gradually decreasing progress if the VDisk is being written to.
- Nonmirrored VDIsks are displayed as a single copy with a progress of 100, and a blank estimated completion time.

The `lsvdisksyncprogress` command also displays the progress of a mirrored VDisk synchronization. After you create a mirrored VDisk using the `svctask mkvdisk` or `svctask addvdiskcopy` command, you can use the command to monitor the progress of the synchronization.

An invocation example

```
svcinfo lsvdisksyncprogress
```

The resulting output

```
| vdisk_id vdisk_name copy_id progress estimated_completion_time  
| 0        vdisk0     0       100  
| 0        vdisk0     1       50      070301150000  
| 3        vdisk3     0       72      070301132225
```

```

| 3      vdisk3    1      100
| 4      vdisk4    0      22      070301160000
| 4      vdisk4    1      100
| 8      vdisk8    0      100
| 8      vdisk8    1      33

```

An invocation example

```
svcinfolsvdisksyncprogress vdisk0
```

The resulting output

```

| vdisk_id vdisk_name copy_id progress estimated_completion_time
| 0        vdisk0    0        100
| 0        vdisk0    1        50      070301150000

```

lsdependentvdisks

Use the **lsdependentvdisks** command to view which volumes will go offline if you remove a specific piece of hardware from the system.

Syntax

```

| ▶▶ svcinfo — lsdependentvdisks — [ -node —node_id_or_name ]
|
| [ -controller —controller_id_or_name_list ] [ -mdisk —mdisk_id_or_name_list ]
|
| [ -drive —drive_id_list ] [ -enclosure —enclosure_id ] [ -canister —canister_id ]
|

```

Parameters

-node

(Optional) Specifies the node for which VDisk dependency is required.

-controller

(Optional) Specifies the controllers for which VDisk dependency is required.

-mdisk

(Optional) Specifies the MDisks for which VDisk dependency is required.

-drive

(Optional) Specifies the drives for which VDisk dependency is required. There is a maximum of 128 entries.

-enclosure

(Optional) Specifies the enclosure for which VDisk dependency is required.

-canister

(Optional) Specifies an enclosure canister if **-enclosure**. This option is not valid for any other type.

Note: Possible values are 1 and 2.

Description

Use this command to view which volumes will go offline if you remove a specific piece of hardware from the system. Use this command before you perform maintenance, to determine which volumes will be affected.

An invocation example

```
svcinfo lsdependentvdisks -delim : -drive 0:1
```

The resulting output

```
vdisk_id:vdisk_name  
4:vdisk4  
5:vdisk5
```

| **Note:** This means that if drives 0 and 1 are removed, then volume vdisk4 and volume vdisk5 will go
| offline.

| lssasfabric

| Use the **lssasfabric** command to see which canisters are visible to a node, and the order of these
| canisters.

| Syntax

| ▶— svcinfo — lssasfabric —————▶

| Description

| Use this command to see which canisters are visible to a node, and the order of these canisters. Table 39
| describes possible outputs.

| *Table 39. lssasfabric output*

Attribute	Description
enclosure_id	The identity of the enclosure the strand goes to.
canister_id	The canister in the enclosure that the strand goes to.
canister_port_id	The canister port that the strand goes to.
control_enclosure_id	The identity of the enclosure the strand comes from.
node_canister_id	The identity of the canister the strand comes from.
node_canister_port_id	The node canister port the strand is from. This should be the same as the chain ID.
position	The position in the strand or chain.
IO_group_id	The I/O group the strand belongs to. This should be the same as the enclosure IO group.
IO_group_name	The I/O group the strand belongs to. This should be the same as the enclosure IO group.
node_id	The identity of the node that the strand is from. This is the same physical object as the node_canister
node_name	The name of the node that the strand is from. This is the same physical object as the node_canister.

| **An invocation example with three enclosures: Enclosure 1 is the control enclosure, Enclosure 2 is on
| chain 1 (node canister port 1) using canister port 1 as its connector, and Enclosure 3 is on chain 2
| (node canister port 2) using canister port 2 as its connector**

```
| svcinfo lssasfabric
```

| The resulting output

| **Note:** In this guide, the following output is split into two parts. This is for illustrative purposes; the output will not appear in two parts when you run this command.

| **Output**

enclosure_id	canister_id	canister_port_id	control_enclosure_id	node_canister_id
1	1	1	1	1
1	2	1	1	2
2	1	1	1	1
2	2	1	1	2
3	1	2	1	1
3	2	2	1	2

| **Output, continued**

node_canister_port_id	position	I0_group_id	I0_group_name	node_id	node_name
2	0	0	io_grp0	1	node1
2	0	0	io_grp0	2	node2
1	1	0	io_grp0	1	node1
1	1	0	io_grp0	2	node2
2	1	0	io_grp0	1	node1
2	1	0	io_grp0	2	node2

showtimezone

Use the **showtimezone** command to display the current time zone settings for the cluster.

Syntax



Parameters

-nohdr

(Optional) By default, headings are displayed for each column of data in a concise style view, and for each item of data in a detailed style view. The **-nohdr** parameter suppresses the display of these headings.

Note: If there is no data to be displayed, headings are not displayed.

-delim delimiter

(Optional) By default in a concise view, all columns of data are space-separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. If you enter **-delim :** on the command line, the colon character (:) separates all items of data in a concise view; for example, the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified character.

Description

This command displays a single time zone and its associated ID. This is the current time zone setting for the cluster. A list of available time-zones can be viewed by running the **svcinfo lstimezones** command. The time zone can be changed by running the **svctask settimezone** command.

An invocation example

```
svcinfo showtimezone -delim :
```

The resulting output

```
id:timezone  
522:UTC
```

resetleds

Use the **resetleds** command to simultaneously switch off all LEDs in the cluster.

Syntax

```
▶▶— svctask —resetleds—▶▶
```

Parameters

None.

Description

The directed maintenance procedures (DMPs) use this command to simultaneously switch off all LEDs in the cluster. This ensures that, when the DMP switches on an identity LED, it is the only one in the cluster. This command works only on LEDs on devices that can be communicated with; it will fail if an object is offline or if the enclosure is an unsupported type. Some LEDs are not affected by this command (for example, LEDs on independently controlled objects, LEDs with hardware-only controls, or LEDs on offline objects).

```
svctask resetleds
```

Chapter 23. Licensing commands

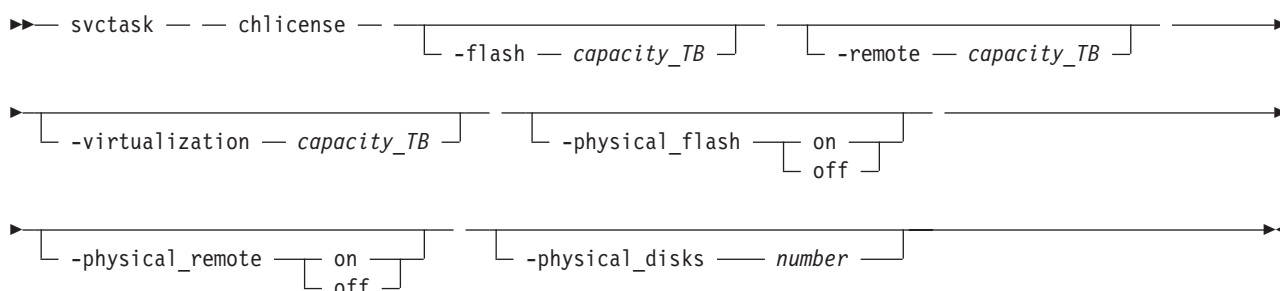
The following commands enable you to work with SAN Volume Controller licensed functions.

- **chlicense**
- **dumpinternallog**

chlicense

The **chlicense** command changes license settings for cluster features.

Syntax



Parameters

-flash *capacity_TB*

(Optional) Changes cluster licensing for the FlashCopy feature. To change the licensed capacity for the FlashCopy feature, specify a capacity in terabytes (TB).

Note: Only use the optional **flash** parameter with the SAN Volume Controller.

| **-remote** *capacity_TB*

| (Optional) Changes cluster licensing for the Metro Mirror and Global Mirror feature. To change the licensed capacity for the Metro Mirror and Global Mirror feature, specify a capacity in terabytes (TB).

| **Note:** For Storwize V7000, specify the total number of internal and external enclosures that you have licensed on your system. You must have a Remote Mirroring license for all enclosures.

| **-virtualization** *capacity_TB*

| (Optional) Changes cluster licensing for the Virtualization feature. To change the licensed capacity for the Virtualization feature, specify a capacity in terabytes (TB).

| **Note:** For Storwize V7000, specify the number of enclosures of external storage that you have been authorized by IBM to use.

| **-physical_flash** **on** | **off**

| (Optional) For physical disk licensing, enables or disables the FlashCopy feature. The default value is **off**.

| **-physical_remote** **on** | **off**

| (Optional) For physical disk licensing, enables or disables the Metro Mirror and Global Mirror feature. The default value is **off**.

-physical_disks *number*

(Optional) Changes the licensed settings of the cluster for physical disk licensing. Enter the number of physical disks that your cluster is licensed to manage. The correct number is identified in your license.

Note:

- If the **-physical_disks** value is set to zero, the **-physical_flash** and **-physical_remote** values are turned off.
- If the **-physical_disks** value is nonzero, the **-flash**, **-remote**, and **-virtualization** values cannot be set.
- If the **-flash**, **-remote**, or **-virtualization** values are nonzero, the **-physical_flash**, **-physical_remote**, and **-physical_disks** values cannot be set.
- If the **-physical_disks** value is nonzero, only the FlashCopy and RemoteCopy usage is monitored and appropriate error messages are logged.

Description

The **chlicense** command changes license settings for the cluster. Any change that is made is logged as an event in the license setting log.

- | For Storwize V7000, the enclosure license already includes virtualization of internal drives on your system. You can use this command to set any additional options. The total amounts for your cluster or clusters must not exceed the total capacity authorization that you have obtained from IBM.
- | For SAN Volume Controller the default is to have no copy services functions licensed, but this does not stop you from creating and using Copy Services. However, errors are placed in the license settings log that state that you are using an unlicensed feature. The command-line tool return code also notifies you that you are using an unlicensed feature.

For Storwize® V7000, the default is to have no Metro Mirror or Global Mirror function licensed, but this does not stop you from creating and using Copy Services. However, errors are placed in the license settings log that state that you are using an unlicensed feature. The command-line tool return code also notifies you that you are using an unlicensed feature.

The total virtualized capacity can also be modified with this command. This is the number of terabytes (TB) of virtual disk capacity that can be configured by the cluster.

When you reach 90% capacity, any attempt to create or extend Virtual Disks, Relationships, or Mappings results in a message from the command-line tool. This does not stop you from creating and expanding Virtual Disks, Relationships, or Mappings. When usage reaches or exceeds 100% capacity, errors are placed in the license settings log.

Any error that is placed in the license settings log results in a generic error being placed in the cluster error log. This occurs when you issue a command that violates the license agreement. The return code also notifies you that you are violating the license settings.

An invocation example

- | `svctask chlicense -remote 5`

The resulting output

No feedback

dumpinternallog

The **dumpinternallog** command dumps the contents of the license settings error and event log to a file on the current configuration node.

Syntax

▶— svctask — — dumpinternallog —————▶

Description

This command dumps the contents of the internal license settings error and event log to a file on the current configuration node.

This file is always called **feature.txt** and is created, or overwritten, in the **/dumps/feature** directory on the configuration node.

This file can be requested by IBM service personnel.

Before making any entries, the license settings log contains only zeros. A dump of this log from the **svctask dumpinternallog** command results in an empty file.

An invocation example

```
svctask dumpinternallog
```

The resulting output

No feedback

Chapter 24. Service mode commands (Discontinued)

Attention: The service mode commands are discontinued.

applysoftware (Discontinued)

Attention: The `svcservicemodetask applysoftware` command is discontinued. Use the `satask installsoftware` command instead.

cleardumps (Discontinued)

Attention: The `svcservicemodetask cleardumps` command is discontinued.

dumperrlog (Discontinued)

Attention: The `svcservicemodetask dumperrlog` command is discontinued.

exit (Discontinued)

Attention: The `svcservicemodetask exit` command is discontinued. Use the `satask stopservice` command instead.

Chapter 25. Service mode information commands (Discontinued)

Attention: The service mode information commands are discontinued.

ls2145dumps (Discontinued)

Attention: The `svcservicemodeinfo ls2145dumps` command is discontinued. Use the `svcinfo lsdumps` command to display a list of files in a particular dumps directory.

lscimomdumps (Discontinued)

Attention: The `svcservicemodeinfo lscimomdumps` command is discontinued. Use the `svcinfo lsdumps` command to display a list of files in a particular dumps directory.

lsclustervpd (Discontinued)

Attention: The `svcservicemodeinfo lsclustervpd` command is discontinued. Use the `sainfo lsservicestatus` command instead.

lserrlogdumps (Discontinued)

Attention: The `svcservicemodeinfo lserrlogdumps` command is discontinued. Use the `svcinfo lsdumps` command to display a list of files in a particular dumps directory.

lsfeaturedumps (Discontinued)

Attention: The `svcservicemodeinfo lsfeaturedumps` command is discontinued. Use the `svcinfo lsdumps` command to display a list of files in a particular dumps directory.

lsiostatsdumps (Discontinued)

Attention: The `svcservicemodeinfo lsiostatsdumps` command is discontinued. Use the `svcinfo lsdumps` command to display a list of files in a particular dumps directory.

lsiotracedumps (Discontinued)

Attention: The `svcservicemodeinfo lsiotracedumps` command is discontinued. Use the `svcinfo lsdumps` command to display a list of files in a particular dumps directory.

lsmdiskdumps (Discontinued)

Attention: The `svcservicemodeinfo lsmdiskdumps` command is discontinued. Use the `svcinfo lsdumps` command to display a list of files in a particular dumps directory.

| **lsnodevpd (Discontinued)**

| **Attention:** The `svcserviceinfo lsnodevpd` command is discontinued. Use the `svcinfo lsnodevpd` command to display the vital product data (VPD) for each node.

| **lssoftwareumps (Discontinued)**

| **Attention:** The `svcserviceinfo lssoftwareumps` command is discontinued. Use the `svcinfo lsumps` command to display a list of files in a particular dumps directory.

Chapter 26. Array commands

Array commands capture information that can assist you with managing arrays.

charray

Use the **charray** command to change array attributes.

Syntax

```
svctask —charray— [-name —new_name_arg—] [-sparegoal —0-100—] [-balanced]
—mdisk_id | mdisk_name—
```

Parameters

-name

(Optional) The new name to apply to the array MDisk.

-sparegoal

(Optional) Sets the number of spares to protect the array members with.

-balanced

(Optional) Forces the array to balance and configure the member spare goals of the present drives.

mdisk_id

Identifies (by ID) which array the MDisk command applies to.

mdisk_name

Identifies (by user-defined name) which array the MDisk command applies to.

Description

This command changes an array's attributes.

Invocation examples

```
svctask charray -name raid6mdisk0 0
svctask charray -sparegoal 2
svctask charray -balanced
```

The resulting output


No feedback

charraymember

Use the **charraymember** command to modify an array member's attributes, or to swap a member of a RAID array with that of another drive.

Syntax

```
svctask —charraymember— -member —member_id—
```



Parameters

-member

Identifies the array member index to be operated on.

-newdrive

(Optional) Identifies the drive you want to add to the array.

-immediate

(Optional) Specifies that the old disk is to be immediately removed from the array, and the new disk rebuilt. If you do not choose this option, exchange is used; this preserves redundancy during the rebuild.

-balanced

(Optional) If specified, forces the array member spare goals to be set to those of present array member, those of an existing exchange, or those of the newDrive.

-unbalanced

(Optional) Forces the array member change if the newDrive does not meet the array member goals.

mdisk_id

(Either the ID or the name is required.) Identifies which array the MDisk command applies to.

mdisk_name

(Either the ID or the name is required.) Identifies which array the MDisk command applies to.

Description

This command modifies an array member's attributes, or to swap a member of a RAID array with that of another drive. Table 40 shows the command combination options.

Table 40. *charraymember combination options*

Option	Description
-balance	<ul style="list-style-type: none"> Member goals are set to the properties of the existing member or exchange drive. The command will fail if the member is not populated with a drive. Member goals are set to the properties of the current member drives being exchanged into the array count as members. If no exchange exists, the existing member drive goals are used.
-newdrive driveId	<ul style="list-style-type: none"> The command processes the exchange, and does NOT update the member goals. You must specify a new drive that is an exact match for the member goals. The command will fail if the drive is not an exact match.
-newdrive driveId -balance	The processes the exchange and updates the member goals to the properties of the new drive.
-newdrive driveId -unbalance	<ul style="list-style-type: none"> The command processes the exchange and does NOT update the member goals. This is only permitted when the array is degraded and the member is empty. This means -immediate is mute, the exchange will always be immediate. Later, if drives are a sufficient member goal match, the array rebalance will select those drives. A balancing exchange will restart the member goals.

| **An invocation example**

| To swap a spare/candidate drive for a member 0 drive using exchange:

```
| svctask charraymember -member 0 -newdrive 4 mdisk2
```

| **An invocation example**

| To swap a spare/candidate drive for a member 1 drive and start component rebuild for the new member:

```
| svctask charraymember -member 1 -newdrive 3 -immediate mdisk3
```

| **An invocation example**

| To swap in a spare/candidate drive for member index 2. If there is an drive present there then the exchange is:

```
| svctask charraymember -member 2 -newdrive 4 mdisk4
```

| **An invocation example**

| To force member 4 to change its spare goals to its associated drive:

```
| svctask charraymember -member 4 -balanced mdisk6
```

| **An invocation example**

| To force an exchange and make the array change its goals to the new drive:

```
| svctask charraymember -member 3 -newdrive 9 -balanced mdisk5
```

| **An invocation example**

| To force an unbalancing exchange when drive 8 does not match the goals:

```
| svctask charraymember -member 2 -newdrive 8 -unbalanced mdisk5
```

| **An invocation example**

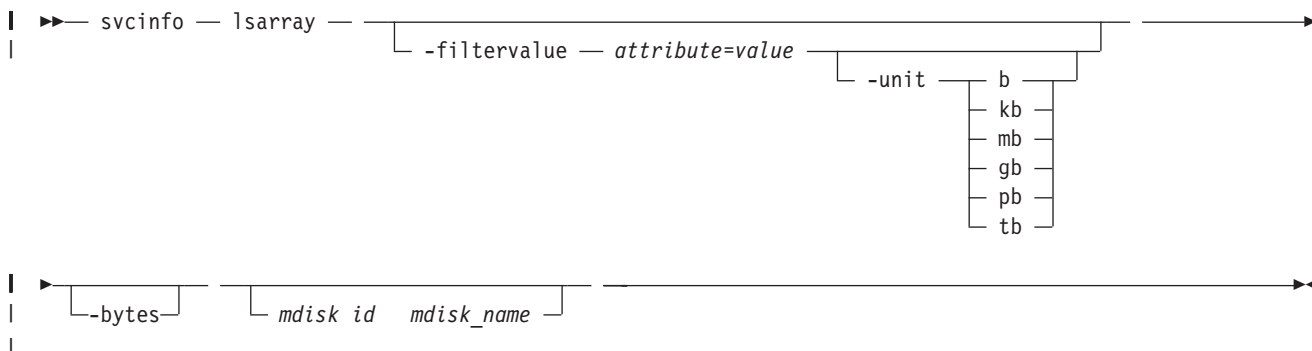
| To force an immediate exchange and make the array change its goals to the new drive:

```
| svctask charraymember -member 3 -newdrive 9 -balanced -immediate mdisk5
```

| **lsarray**

| Use the **lsarray** command to list the array MDisks.

| **Syntax**



| Parameters

| **-filtervalue** *attribute=value*

| (Optional) Specifies a list of one or more filter attributes matching the specified values; see **-filtervalue?** for the supported attributes. Only objects with a value that matches the filter attribute value are returned. If **capacity** is specified, the units must also be included.

| **Note:** Some filters allow the use of a wildcard when entering the command. The following rules apply to the use of wildcards with the SAN Volume Controller CLI:

- | • The wildcard character is an asterisk (*).
- | • The command can contain a maximum of one wildcard, which must be the first or last character in the string.
- | • When using a wildcard character, you must enclose the filter entry within double quotation marks (""), as follows:

| `svcinfolarray -filtervalue "name=md*"`

| **-filtervalue?**

| (Optional) Includes all of the valid filter attributes in the report. The following filter attributes are valid for the **svcinfolarray** command:

- | • `mdisk_id`
- | • `mdisk_name`
- | • `status`
- | • `mode`
- | • `mdisk_grp_id`
- | • `mdisk_grp_name`
- | • `capacity`
- | • `fast_write_state`
- | • `raid_status`
- | • `raid_level`
- | • `redundancy`
- | • `strip_size`
- | • `write_verify`
- | • `spare_goal`
- | • `spare_protection_min`
- | • `balanced`
- | • `tier`

| Any parameters specified with the **-filtervalue?** parameter are ignored.

| **-bytes**

| (Optional) Requests output of capacities in bytes (instead of rounded values).

| **mdisk_id**

| (Optional) The identity of the array MDisk.

| **mdisk_name**

| (Optional) The MDisk name that you provided.

| Description

| This command returns a concise list or a detailed view of array MDisks visible to the cluster. The **lsmdisk** command provides the potential output for array MDisks.

Table 41. MDisk output

Attribute	Values
status	<ul style="list-style-type: none"> • online • offline • excluded • degraded (applies only to internal MDisks)
mode	unmanaged, managed, image, array
quorum_index	0, 1, 2, or blank if the MDisk is not being used as a quorum disk
block_size	512, 524 bytes in each block of storage
ctrl_type	4, 6, where 6 is a solid-state drive (SSD) attached inside a node and 4 is any other device
tier	<p>The tier this MDisk has been assigned to by auto-detection (for internal arrays) or by the user:</p> <ul style="list-style-type: none"> • generic_ssd • generic_hdd (the default value for newly discovered or external MDisk) <p>Note: You can change this value using the <code>chmdisk</code> command.</p>
raid_status	<ul style="list-style-type: none"> • offline - the array is offline on all nodes • degraded - the array has deconfigured or offline members; the array is not fully redundant • syncing - array members are all online, the array is syncing parity or mirrors to achieve redundancy • initting - array members are all online, the array is initializing; the array is fully redundant • online - array members are all online, and the array is fully redundant
raid_level	The RAID level of the array (RAID0, RAID1, RAID5, RAID6, RAID10).
redundancy	The number of how many member disks can fail before the array fails.
strip_size	The strip size of the array (in KB).
spare_goal	The number of spares that the array members should be protected by.
spare_protection_min	The minimum number of spares that an array member is protected by.
balanced	<p>Describes if the array is balanced to its spare goals:</p> <ul style="list-style-type: none"> • exact: all populated members have exact capability match, exact location match • yes: all populated members have at least exact capability match, exact chain, or different enclosure or slot • no: anything else

The following define the status fields:

Online

The MDisk is online and available.

Degraded

(Internal MDisks only) The array has members that are degraded, or the `raid_status` is degraded.

Offline

All paths to the MDisk are lost.

Excluded

The MDisk is excluded from use by the cluster; the MDisk port error count exceeded the threshold.

| A concise invocation example

```
| svcinfo lsarray -delim :
```

| The resulting output

```
| mdisk_id:mdisk_name:status:mdisk_grp_id:mdisk_grp_name:capacity:raid_status:
| raid_level:redundancy:strip_size:tier
| 1:mdisk1:online:0:mdiskgrp0:68.4GB:online:raid0:0:256:generic_hdd
| 2:mdisk2:online:0:mdiskgrp0:88.4GB:syncing:raid5:1:256:generic_hdd
| 533:mdisk533:degraded:1:mdiskgrp1:78.2GB:syncing:raid6:2:128:generic_hdd
| 534:mdisk534:online:2:mdiskgrp1:94.2GB:initting:raid6:2:64:generic_ssd
```

| A full invocation example

```
| svcinfo lsarray mdisk1
```

| The resulting output

```
| mdisk_id:1
| mdisk_name:mdisk1
| status:online
| mode:array
| mdisk_grp_id:0
| mdisk_grp_name:mdiskgrp0
| capacity:68.4GB
| quorum_index:
| block_size:
| controller_name:
| ctrl_type:
| ctrl_WWNN:
| controller_id:
| path_count:
| max_path_count:
| ctrl_LUN_#:
| UID:
| preferred_WWPN:
| active_WWPN:
| node_id:
| node_name:
| location:
| fast_write_state:empty
| raid_status:online
| raid_level:raid0
| redundancy:0
| strip_size:256
| spare_goal:2
| spare_protection_min:2
| balanced:yes
| tier:generic_hdd
```

| **Isarrayinitprogress**

| Use the **Isarrayinitprogress** command to view the progress of array background initialization that occurs after creation.

| **Syntax**

```
| ►► svcinfo — lsarrayinitprogress —————┬───────────────────┘►►
|                                     └mdisk id | mdisk_name┘
```

| **Parameters**

| **mdisk_id**

| (Optional) The identity of the array MDisk.

| **mdisk_name**
| (Optional) The user-defined MDisk name.

| **Description**

| This command shows the progress of array background initialization. Table 42 shows possible outputs.

| *Table 42. Isarrayinitprogress output*

Attribute	Value
progress	The percentage of initialization task that has been completed.
estimated_completion_time	The expected initialization task completion time (YYMMDDHHMMSS).

| **A concise invocation example**

| `svcinfolarrayinitprogress -delim :`

| **The resulting output**

| `mdisk_id:mdisk_name:progress:estimated_completion_time`
| `0:mdisk0:50:070301120000`
| `1:mdisk1:51:070301130000`
| `2:mdisk2:32:070301153500`

| **A concise invocation (qualified with MDisk) example**

| `svcinfolarrayinitprogress -delim : mdisk2`

| **The resulting output**

| `mdisk_id:mdisk_name:progress:estimated_completion_time`
| `2:mdisk2:32:070301153500`

| **An invocation example for an array that has finished initialization**

| `svcinfolarrayinitprogress -delim : mdisk4`

| **The resulting output**

| `mdisk_id:mdisk_name:progress:estimated_completion_time`
| `4:mdisk4:100:`

| **Isarraylba**

| Use the **Isarraylba** command to permit an array logical block address (LBA) to be found from a drive and LBA.

| **Syntax**

| `▶▶— svcinfolarraylba — [-drive1lba — lba — -drive — drive_id]`

| **Parameters**

| **-drive1lba**

| The LBA on drive that you want to convert to the array LBA.

| **-drive**

| The ID of the drive you want to view.

Description

This command enables you to permit an array LBA to be found from a drive and LBA. Table 43 shows possible outputs.

Table 43. *lsarraylba* output

Attribute	Value
type	The type of MDisk extent allocation: <ul style="list-style-type: none">• allocated• unallocated
mdisk_lba	The LBA on the array MDisk (blank if none).
mdisk_start	The start of range of LBAs (strip) on the array MDisk (blank if none).
mdisk_end	The end of range of LBAs (strip) on the array MDisk (blank if none).
drive_start	The start of range of LBAs (strip) on the drive (blank if none).
drive_end	The end of range of LBAs (strip) on the drive (blank if none).

The following example shows:

- How a drive 2 LBA -*xff* maps to MDisk 2 LBA 0*xff*
- How the first 255 LBAs of drive 2 map to the second 255 LBAs of the MDisk.

An invocation example

```
svcinfolarraylba -drive1lba 0xff -drive 2 -delim :
```

The resulting output

```
mdisk_id:mdisk_name:type:mdisk_lba:mdisk_start:mdisk_end:drive_start:drive_end
0:mdisk2:allocated:0x000000000000001ff:0x0000000000000100:0x0000000000001ff:
0x0000000000000000:0x00000000000000ff
```

lsarraymember

Use the *lsarraymember* command to list the member drives of one or more array MDisks.

Syntax

```
svcinfolarraymember [-bytes—mdisk_id—mdisk_name]
```

Parameters

mdisk_id
(Optional) The identity of the array MDisk.

mdisk_name
(Optional) The MDisk name that you provided.

Description

This command lists the member drives of one or more array MDisks. It describes positions within an array that not occupied by a drive. The positions determine how mirroring the RAID's will take place; for example, it determines if *x* is mirrored to *y* for RAID-10, where parity starts from RAID-5.

Table 44 on page 391 provides the potential output for this command.

Table 44. *lsarraymember* output

Attribute	Value
member_id	The identity of the array member; represents drive order in RAID array
drive_id	The identity of the drive for member ID, or the source drive if an exchange is in progress (blank if there is none configured).
new_drive_id	The ID of the drive being exchanged with this member ID (blank if there is none)
spare_protection	The number of spares protecting the array member
balanced	If the array member drive matches the spare goals: <ul style="list-style-type: none"> exact - exact capability match, exact location match yes - exact capability match, exact chain, different enclosure or slot no - anything else (blank) - there is no drive configured for the member

A concise invocation example

```
svcinfolarraymember -delim :
```

The resulting output

```
svcinfolarraymember -delim :
mdisk_id:mdisk_name:member_id:drive_id:new_drive_id:spare_protection:balanced
2:mdisk1:0:55::1:exact
2:mdisk1:1:56::1:exact
2:mdisk2:0:0::2:exact
2:mdisk2:1:2:5:3:exact
2:mdisk2:2:::
2:mdisk2:3:8::0:no
```

A concise invocation example (qualified with MDisk)

```
svcinfolarraymember mdisk2 -delim :
```

The resulting output

```
mdisk_id:mdisk_name:member_id:drive_id:new_drive_id:spare_protection:balanced
2:mdisk2:0:0::2:exact
2:mdisk2:1:2:5:3:exact
2:mdisk2:2:::
2:mdisk2:3:8::0:no
```

Note: From this output, you can see that:

- The array has four members (possibly a 4-member RAID-10 array).
- The second array member is undergoing exchange for drive5.
- The third array member is not configured. It might have gone offline or failed, without a hot spare available.
- The fourth array member has no spare protection and is not balanced.

An invocation example (two arrays)

```
svcinfolarraymember -delim :
```

The resulting output

```
mdisk_id:mdisk_name:member_id:drive_id:new_drive_id:spare_protection:balanced
2:mdisk1:0:55::1:exact
2:mdisk1:1:56::1:exact
```

```
| 2:mdisk2:0:0:::2:exact
| 2:mdisk2:1:2:5:::3:exact
| 2:mdisk2:2:::0:
| 2:mdisk2:3:8:::0:no
```

| **An invocation example (an array expanding from membership (55,56) to (55,57,58))**

```
| svcinfo lsarraymember -delim : mdisk3
```

| **The resulting output**

```
| mdisk_id:mdisk_name:member_id:drive_id:new_drive_id:spare_protection:balanced
| 3:mdisk3:0:55::55:1:exact
| 3:mdisk3:1:56::57:1:exact
| 3:mdisk3:2:::58:1:exact
```

| **An invocation example (an array contracting from membership (55,57,58) to (55,56))**

```
| svcinfo lsarraymember -delim : mdisk3
```

| **The resulting output**

```
| mdisk_id:mdisk_name:member_id:drive_id:new_drive_id:spare_protection:balanced
| 3:mdisk3:0:55::55:1:exact
| 3:mdisk3:1:57::56:1:exact
| 3:mdisk3:2:58:::1:exact
```

| **Isarraymembergoals**

| Use the **Isarraymembergoals** command to list the spare goals for member drives of one or more array MDisks.

| **Syntax**

```
| ▶▶— svcinfo — lsarraymembergoals —————▶▶
|                                     └─bytes—mdisk id—mdisk_name—┘
```

| **Parameters**

- | **-bytes**
| (Optional) Requests output of capacities in bytes (instead of rounded values).
- | **mdisk_id**
| (Optional) The identity of the array MDisk.
- | **mdisk_name**
| (Optional) The MDisk name that you provided.

| **Description**

| This command list the spare goals for member drives of one or more array MDisks. Table 45 provides the potential output for this command.

| *Table 45. Isarraymembergoals output*

Attribute	Values
member_id	The ID of the array member which represents the drive order in the RAID array.
drive_id	The ID of the drive for the member ID (blank if none is configured).
capacity_goal	The capacity goal for the array member (same for all members in the array).

Table 45. *lsarraymembergoals* output (continued)

Attribute	Values
tech_type_goal	The technology goal for the array member: <ul style="list-style-type: none"> • sas_ssd • sas_hdd • sas_nearline_hdd
RPM_goal	The RPM goal for array member (blank for SSDs).
enclosure_id_goal	The ID of the member enclosure goal (blank if any can be selected).
slot_id_goal	The ID of the member slot goal.

An invocation example (a four-member RAID 10 SAS array that is split across chains)

```
svcinfolarraymembergoals mdisk2 -delim :
```

The resulting output

```
mdisk_id:mdisk_name:member_id:drive_id:capacity_goal:
tech_type_goal:RPM_goal:enclosure_id_goal:slot_id_goal
2:mdisk2:0:0:68.4GB:sas_hdd:15000:1:1
2:mdisk2:1:17:68.4GB:sas_hdd:15000:1:2
2:mdisk2:2:1:68.4GB:sas_hdd:15000:14:1
2:mdisk2:3:18:68.4GB:sas_hdd:15000:14:2
```

An invocation example (a six-member RAID 10 SAS or SATA array)

```
svcinfolarraymembergoals mdisk3 -delim :
```

The resulting output

```
mdisk_id:mdisk_name:member_id:drive_id:capacity_goal:
tech_type_goal:RPM_goal:enclosure_id_goal:slot_id_goal
3:mdisk3:0:10:155.0GB:sas_ssd::1:4
3:mdisk3:1:21:155.0GB:sas_hdd:15000:2:3
3:mdisk3:2:12:155.0GB:sas_nearline_hdd:7200:7:3
3:mdisk3:4:23:155.0GB:sas_ssd::2:2
3:mdisk3:5:14:155.0GB:sas_nearline_hdd:7200:9:3
3:mdisk3:6:25:155.0GB:sas_hdd:15000:2:8
```

An invocation example (a four-member RAID 0 SAS array contained within a single enclosure)

```
svcinfolarraymembergoals mdisk4 -delim :
```

The resulting output

```
mdisk_id:mdisk_name:member_id:drive_id:capacity_goal:
tech_type_goal:RPM_goal:enclosure_id_goal:slot_id_goal
2:mdisk2:0:0:222.0GB:sas_nearline_hdd:15000:1:1
2:mdisk2:1:1:222.0GB:sas_nearline_hdd:15000:1:2
2:mdisk2:2:2:222.0GB:sas_nearline_hdd:15000:1:3
2:mdisk2:3:3:222.0GB:sas_nearline_hdd:15000:1:4
```

lsarraymemberprogress

Use the `lsarraymemberprogress` command to display array member background process status.

Syntax

```
svcinfolarraymemberprogress mdisk_id—mdisk_name
```

| **Parameters**

| **mdisk_id**
| (Optional) The identity of the array MDisk.

| **mdisk_name**
| (Optional) The MDisk name that you provided.

| **Description**

| This command displays array member background process status. Exchange cannot start on a rebuilding member because both component rebuild and exchange are shown in the same view. Table 46 provides the potential output for this command.

| *Table 46. Isarraymemberprogress output*

Attribute	Value
member_id	The array member index.
drive_id	The ID of the drive.
task	The identity of task: <ul style="list-style-type: none">• rebuild• exchange
new_drive_id	The identity of drive being exchanged.
progress	The task percentage complete.
estimated_completion_time	The expected task completion time (YYMMDDHHMMSS; blank if completion time unknown).

| **A concise invocation example**

| `svcinfolarraymemberprogress -delim :`

| **The resulting output**

| `mdisk_id:mdisk_name:member_id:drive_id:task:new_drive_id:progress:estimated_completion_time`
| `0:mdisk0:2:3:rebuild::50:070301120000`
| `1:mdisk1:0:5:rebuild::51:070301130000`
| `2:mdisk2:4:1:exchange:12:32:070301153500`
| `2:mdisk2:5:16:exchange:13:0:`
| `2:mdisk2:5:17:exchange:14:0:`

| **An MDisk qualified concise example**

| `svcinfolarraymemberprogress mdisk2`

| **The resulting output**

| `mdisk_id:mdisk_name:member_id:drive_id:task:new_drive_id:progress:estimated_completion_time`
| `2:mdisk2:4:1:exchange:12:32:070301153500`
| `2:mdisk2:5:16:exchange:13:0:`
| `2:mdisk2:5:17:exchange:14:0:`

| **Isarraysyncprogress**

| The `Isarraysyncprogress` command displays how synchronized a RAID array is.

Syntax

```
svcinfo --lsarraysyncprogress --disk_id --disk_name
```

Parameters

mdisk_id
(Optional) The ID of the MDisk you want to view.

mdisk_name
(Optional) The user-defined name of the MDisk you want to view.

Description

This command shows you how synchronized a RAID array is. It includes internal activity that is working toward a fully synchronized array. Table 47 provides the potential output.

Table 47. *lsarraysyncprogress* output

Attribute	Value
progress	The percentage of the array that is synchronized.
estimated_completion_time	The expected synchronization completion time (YYMMDDHHMMSS; blank if completion time unknown).

A concise invocation example

```
svcinfo lsarraysyncprogress -delim :
```

The resulting output

```
mdisk_id:mdisk_name:progress:estimated_completion_time
0:mdisk0:50:070301120000
1:mdisk1:51:070301130000
2:mdisk2:32:070301153500
```

A concise view (qualified with mdisk id for mdisk2) invocation example

```
svcinfo lsarraysyncprogress -delim : mdisk2
```

The resulting output

```
mdisk_id:mdisk_name:progress:estimated_completion_time
2:mdisk2:32:070301153500
```

A concise view (qualified with mdisk id for in sync mdisk10) invocation example

```
svcinfo lsarraysyncprogress -delim : mdisk10
```

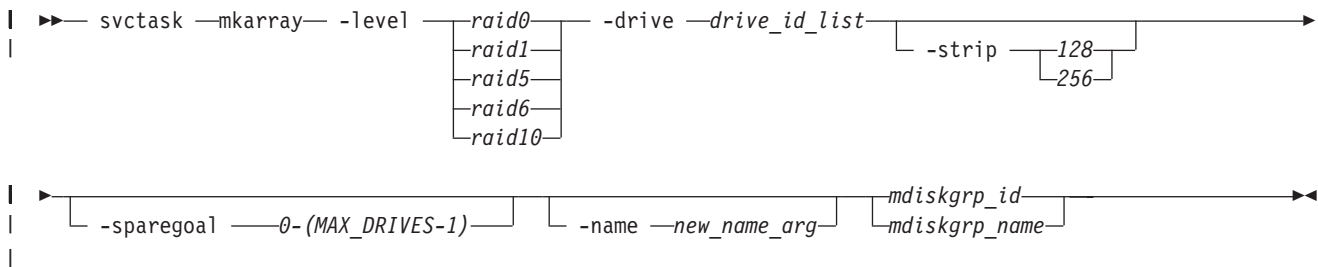
The resulting output

```
mdisk_id:mdisk_name:progress:estimated_completion_time
0:mdisk0:100:
```

mkarray

Use the **mkarray** command to create an MDisk RAID array and add it to an MDisk group.

Syntax



Parameters

-level

Sets the RAID level for the array MDisk being created.

-drive

Identifies the drive or drives to use as members of the RAID array.

-strip

(Optional) Sets strip size (in kilobytes) for the array MDisk being created.

-sparegoal

(Optional) Sets the number of spares that this array's members should be protected by.

-name

(Optional) Specifies the name to which you want to apply the array MDisk.

mdiskgrp_id

Identifies the MDisk group (by ID) to which you want to add the created array mdisk.

mdiskgrp_name

Identifies MDisk group (by the user-defined name) to which you want to add created array MDisks.

Description

This command creates an array MDisk RAID array and adds it to an MDisk group. Although the array's tier is automatically determined, you can change it later using the **chmdisk** command.

Standard output

MDisk, id [x], successfully created

An invocation example (to create fully redundant arrays)

```
svctask mkarray -level raid0 -drive 0:1:2:3 raid0grp
```

The resulting output

MDisk, id [0], successfully created

An invocation example (to create fully redundant arrays)

```
svctask mkarray -level raid1 -drive 4:5 -strip 128 mdiskgrp4
```

The resulting output

MDisk, id [1], successfully created

An invocation example (to create fully redundant arrays)

```
svctask mkarray -level raid6 -drive 6:7:8:9:10 raid6grp
```

The resulting output

| MDisk, id [2], successfully created

| **recoverarray**

| Use the **recoverarray** command to recover a specific corrupt array in a dead domain scenario.

| **Syntax**

| ►►— `svctask —recoverarray—mdisk_id—mdisk_name`—►►

| **Parameters**

| **mdisk_id**

| Identifies (by ID)

| **mdisk_name**

| Identifies (by user-assigned name)

| **Description**

| This command recovers a specific corrupt array. An array has metadata that represents ongoing/pending platform writes that is lost when the domain nodes are lost.

| **An invocation example**

| `svctask recoverarray mdisk1`

| **The resulting output**

| There is no output if the command is successful.

| **recoverarraybycluster**

| Use the **recoverarraybycluster** command to recover a specific corrupt array in a dead domain scenario.

| **Syntax**

| ►►— `svctask —recoverarraybycluster`—►►

| **Parameters**

| None.

| **Description**

| Use this command to recover the corrupt arrays in a dead domain scenario.

| **An invocation example**

| `svctask recoverarraybycluster`

| **The resulting output**

| There is no output if the command is successful.

| **rmarray**

| Use the **rmarray** command to remove an array mdisk from the configuration.

| **Syntax**

```
| ▶▶— svctask —rmarray— -mdisk —mdisk_id_list—mdisk_name_list—————▶▶  
|  
| ▶ |—————▶▶  
| | —force —mdisk_group_id—mdisk_group_name— |
```

| **Parameters**

| **-mdisk**

| Identifies the array MDisk or MDisks to remove from the mdisk group.

| **-force**

| (Optional) Forces a remove when the MDisk has allocated extents by migrating the used extents to free extents in the mdiskgrp.

| **mdiskgrp_id**

| Identifies (by ID) the MDisk group to add the created array MDisk to.

| **mdiskgrp_name**

| Identifies (by user-defined name) the MDisk group to add the created array MDisk to.

| **Description**

| This command removes an array mdisk from the configuration. Each array is divided into candidate drives.

| **An invocation example**

```
| svctask rmarray -mdisk 6 mdiskgrp10
```

| **The resulting output**

| There is no output if the command is successful.

Chapter 27. Service information commands

Service information commands are used to view the status of the hardware and report hardware errors.

chnodeled

Use the **chnodeled** command to turn on or off the location LED for the specified canister.

Syntax

```
▶▶ satask -- chnodeled --on | -off--panel_name
```

Parameters

on | *off*

Turns on or off the location LED for the specified canister.

Description

This command turns on or off the canister location LED.

Note: The location LED is mapped onto the physical LEDs using different methods, depending on your hardware. Refer to the documentation for your hardware platform for more information.

To turn on the LED for canister 1 in enclosure 02:

```
satask chnodeled -on 02-1
```

lscmdstatus

Use the **lscmdstatus** command to display the status of any currently running service-aid task.

Syntax

```
▶▶ sainfo -- lscmdstatus --panel_name
```

Parameters

panel_name

The name of the panel. This command will fail if the *panel_name* ID is not in the list returned by **lsservicenodes**.

Description

This command displays the status of any currently running service-aid task. If no task is running, then the completion status of the last task will be displayed.

If no service-aid tasks have run since the node was last restarted, the command will return immediately with no output. Otherwise, it will display something similar to the following:

```
Backup date 20100706 15:53 : quorum time 20100706 16:24
```

lsfiles

Use the **lsfiles** command to display the files on the node that you want to retrieve with the **satask cpfiles** command.

Syntax

```
▶▶ sainfo — lsfiles [-prefix] [path] [panel_name] ▶▶
```

Parameters

panel_name

The name of the panel. The command will fail if the *panel_name* ID is not in the list returned by the **lsservicenodes** command.

-prefix*path*

The path must exist in a permitted listable directory. You can use the following **-prefix** paths:

- /dumps (the default if **-prefix** is not set)
- /dumps/audit
- /dumps/cimom
- /dumps/configs
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/mdisk
- /dumps/syslogs
- /home/admin/upgrade

Description

This command displays a list of the files on the node that you want to retrieve using the **satask cpfiles** command.

```
filename  
sublun.trc.old  
sublun.trc  
100050.trc.old  
eccore.100050.100305.183051  
eccore.100050.100305.183052  
ethernet.100050.trc  
100050.trc
```

lsservicenodes

Use the **lsservicenodes** command to displays a list of all the nodes that can be serviced using the service assistant CLI.

Syntax

```
▶▶ sainfo — — lsservicenodes ▶▶
```

Parameters

None

Description

This command displays a list of all the nodes that can be serviced using the service assistant CLI. This list includes:

- The local node
- Any other nodes in the same cluster as the local node that the local node can communicate with

Nodes not clustered with the local node will not be shown unless they are the partner node. Table 48 shows possible outputs.

Table 48. *Isservicenodes* outputs

Attribute	Value
panel_name	The front panel name, enclosure IDs, or canister IDs that identify the node.
cluster_id	Blank if node is a candidate; otherwise, the value is determined from vpd_cluster.
cluster_name	Blank if node is a candidate; otherwise, the value is determined from vpd_cluster.
node_id	Blank if node is a candidate; otherwise, the value is determined from vpd_cluster.
node_name	Blank if node is a candidate; otherwise, the value is determined from vpd_cluster.
relation	<ul style="list-style-type: none">• Local: the node the CLI command was issued from.• Partner: the node in the same enclosure as the local node.• Cluster: nodes other than the partner that are in the same cluster as the local node.
node_status	<ul style="list-style-type: none">• Active: the node is part of a cluster and can perform I/O.• Service: the node is in service, standby, or node rescue.• Candidate: the node is not part of a cluster or in service.• Starting: the node is part of a cluster and is attempting to join the cluster, and cannot perform I/O.
error_data	Outstanding error and error data, by priority.

Storwize® V7000:

This command displays a list of all the nodes that can be serviced using the service assistant CLI. This list includes:

- The local node
- The partner node in the enclosure
- Any other nodes in the same cluster as the local node that the local node can communicate with

Nodes not clustered with the local node will not be shown unless they are the partner node. Table 49 shows possible outputs.

Table 49. *Isservicenodes* outputs

Attribute	Value
panel_name	The front panel name, enclosure IDs, or canister IDs that identify the node.
cluster_id	Blank if node is a candidate; otherwise, the value is determined from vpd_cluster.
cluster_name	Blank if node is a candidate; otherwise, the value is determined from vpd_cluster.
node_id	Blank if node is a candidate; otherwise, the value is determined from vpd_cluster.
node_name	Blank if node is a candidate; otherwise, the value is determined from vpd_cluster.

Table 49. *Isservicenodes outputs (continued)*

Attribute	Value
relation	<ul style="list-style-type: none"> Local: the node the CLI command was issued from. Partner: the node in the same enclosure as the local node. Cluster: nodes other than the partner that are in the same cluster as the local node.
node_status	<ul style="list-style-type: none"> Active: the node is part of a cluster and can perform I/O. Service: the node is in service, standby, or node rescue. Candidate: the node is not part of a cluster or in service. Starting: the node is part of a cluster and is attempting to join the cluster, and cannot perform I/O.
error_data	Outstanding error and error data, by priority.

```
panel_name:cluster_id:cluster_name:node_id:node_name:relation:node_status:error
07-1:::::local:candidate:
07-2:::::partner:candidate:578
```

Isservicerecommendation

Use the **Isservicerecommendation** command to determine what actions should be performed when servicing a node.

Syntax

```
▶▶ sainfo -- Isservicerecommendation panel_name
```

Parameters

panel_name

(Optional) If no panel ID is provided, the service recommendation for the local node is returned. If a *panel_name* from the list returned by **Isservicenodes** is specified, then the service recommendation for that node is returned. The command will fail if the *panel_name* is not in the list returned by **Isservicenodes**.

Description

This command enables you to determine what actions should be performed when servicing a node.

Example 1 for service_action:

Use fabric tools to diagnose and correct fibre channel fabric problem.

Example 2 for service_action

No service action required, use console to manage node.

Isservicestatus

Use the **Isservicestatus** command to display the current status of a node.

Syntax

```
▶▶ sainfo -- Isservicestatus panel_name
```

Parameters

panel_name

If a *panel_name* is provided, the service recommendation for the local node is returned. If a *panel_name* from the list returned by **Isservicenodes** is specified, then the service recommendation for that node is returned. The command will fail if the *panel_name* ID is not in the list returned by **Isservicenodes**. This output will also be returned as the node status on all USB flash drive commands.

Description

Use this command to display the current status of a node. This command provides all the information that can be obtained using the front panel of a SAN Volume Controller node. You can run this command on any node, even one that is not part of a cluster, to obtain the VPD and error status.

Table 50 shows possible outputs.

Table 50. *Isservicestatus* output

Attribute	Value
panel_name	The front panel name, enclosure IDs, or canister IDs that identify the node.
cluster_id	Specifies the ID of a cluster.
cluster_name	Specifies the name of a cluster. When you use this parameter, the detailed view of the specific cluster is displayed and any value that you specified by the -filtervalue parameter is ignored. If you do not specify the <i>cluster_name</i> parameter, the concise view of all clusters that match the filtering requirements that are specified by the -filtervalue parameter are displayed.
cluster_status	The error code is the same as the one displayed on the front panel.
cluster_ip_count	The maximum number of management addresses you can configure.
cluster_ip_port	This, and fields down to prefix_6, are repeated for each management address.
cluster_ip	The IPv4 management IP address.
cluster_gw	The IPv4 management IP gateway.
cluster_mask	The IPv4 management IP mask.
cluster_ip_6	The IPv6 management IP address.
cluster_gw_6	The IPv6 management IP gateway.
cluster_prefix_6	The IPv6 management IP prefix.
node_id	The ID of the node that is being configured.
node_name	The name of the node that is being configured.
node_status	active starting service candidate
config_node	yes no
hardware	8F2 8F4 8G4 CF8 8A4 other
service_IP_address	The IPv4 service address for the node.
service_gateway	The IPv4 service gateway for the node.
service_subnet_mask	The IPv4 service mask for the node.
service_IP_address_6	The IPv6 service address for the node.
service_gateway_6	The IPv6 service gateway for the node.
service_prefix_6	The IPv6 service gateway for the node.
node_sw_version	The software version of the node.
node_sw_build	The build string for software on the node.
cluster_sw_build	The CSM build that the cluster is running.

| *Table 50. Isservicestatus output (continued)*

Attribute	Value
node_error_count	The number of node errors.
node_error_data	The type of node errors.
FC_port_count	The number of FlashCopy ports.
FC_port_id	The beginning of repeating fields for each FlashCopy port; the whole set of fields indicated is repeated for each port.
port_status	This should match the front panel.
port_speed	This should match the front panel.
port_WWPN	The worldwide port number of the port.
SFP_type	long-wave short-wave
ethernet_port_count	The number of detected Ethernet ports.
ethernet_port_id	Specifies the ID of an Ethernet port.
port_status	online offline not configured
port_speed	10Mb/s 100Mb/s 1Gb/s 10Gb/s full half
MAC	A single MAC address.
product_mtm	The machine type and model.
product_serial	The node serial number.
time_to_charge	The estimated start time (in minutes) needed for 50% of the battery to be charged.
battery_charging	The percentage of charge of the batteries.
disk_WWNN_prefix	The most recently used WWNN prefix.
node_WWNN	N/A
enclosure_WWNN_1	N/A
enclosure_WWNN_2	N/A
node_part_identity	N/A
node_FRU_part	N/A
enclosure_part_identity	N/A
PSU_count	N/A
PSU_id	N/A
PSU_status	N/A
battery_count	N/A
battery_id	N/A
battery_status	N/A

| **Storwize® V7000:** Table 51 shows possible outputs.

| **Note:** On a node that is not part of a cluster, some of the fields will be blank or N/A.

| *Table 51. Isservicestatus output*

Attribute	Value
panel_name	The front panel name, enclosure IDs, or canister IDs that identify the node.
cluster_id	Specifies the ID of a cluster.

Table 51. *Isservicestatus* output (continued)

Attribute	Value
cluster_name	Specifies the name of a cluster. When you use this parameter, the detailed view of the specific cluster is displayed and any value that you specified by the -filtervalue parameter is ignored. If you do not specify the <i>cluster_name</i> parameter, the concise view of all clusters that match the filtering requirements that are specified by the -filtervalue parameter are displayed.
cluster_status	The error code is the same as the one displayed on the front panel.
cluster_ip_count	The maximum number of management addresses you can configure.
cluster_ip_port	This, and fields down to <i>prefix_6</i> , are repeated for each management address.
cluster_ip	The IPv4 management IP address.
cluster_gw	The IPv4 management IP gateway.
cluster_mask	The IPv4 management IP mask.
cluster_ip_6	The IPv6 management IP address.
cluster_gw_6	The IPv6 management IP gateway.
cluster_prefix_6	The IPv6 management IP prefix.
node_id	The ID of the node that is being configured.
node_name	The name of the node that is being configured.
node_status	active starting service candidate
config_node	yes no
hardware	8F2 8F4 8G4 CF8 8A4 other
service_IP_address	The IPv4 service address for the node.
service_gateway	The IPv4 service gateway for the node.
service_subnet_mask	The IPv4 service mask for the node.
service_IP_address_6	The IPv6 service address for the node.
service_gateway_6	The IPv6 service gateway for the node.
service_prefix_6	The IPv6 service gateway for the node.
node_sw_version	The software version of the node.
node_sw_build	The build string for software on the node.
cluster_sw_build	The CSM build that the cluster is running.
node_error_count	The number of node errors.
node_error_data	The type of node errors.
FC_port_count	The number of FlashCopy ports.
FC_port_id	The beginning of repeating fields for each FlashCopy port; the whole set of fields indicated is repeated for each port.
port_status	This should match the front panel.
port_speed	This should match the front panel.
port_WWPN	The worldwide port number of the port.
SFP_type	long-wave short-wave
ethernet_port_count	The number of detected Ethernet ports.
ethernet_port_id	Specifies the ID of an Ethernet port.
port_status	online offline not configured
port_speed	10Mb/s 100Mb/s 1Gb/s 10Gb/s full half

| *Table 51. Isservicestatus output (continued)*

Attribute	Value
MAC	A single MAC address.
product_mtm	The machine type and model.
product_serial	The node serial number.
time_to_charge	The estimated start time (in minutes) needed for 50% of the battery to be charged.
battery_charging	The percentage of charge of the batteries.
node_WWNN	The last active WWNN stored in the node; blank if no cluster data.
enclosure_WWNN_1	Canister 1 WWNN from the enclosure VPD.
enclosure_WWNN_2	Canister 2 WWNN from the enclosure VPD.
node_part_identity	The 11S string from the hardware VPD.
node_FRU_part	if stored in node VPD
enclosure_part_identity	The S11 data.
PSU_count	The number of expected PSUs (two).
PSU_id	The ID of the slot the PSU is in.
PSU_status	missing failed active
battery_count	The number of expected batteries (two).
battery_id	The ID of the slot the battery is in.
battery_status	missing failed charging active
node_location_copy	Equivalent to the panel name; blank if a node has been removed from a cluster.
node_product_mtm_copy	Equivalent to panel product_mtm; blank if a node has been removed from a cluster.
node_product_serial_copy	Equivalent to product_serial; blank if a node has been removed from a cluster.
node_WWNN_1_copy	Equivalent to enclosure_WWNN_1; blank if a node has been removed from a cluster.
node_WWNN_2_copy	Equivalent to enclosure_WWNN_2; blank if a node has been removed from a cluster.
latest_cluster_id	The cluster ID running on the current enclosure; blank if a node has been removed from a cluster.
next_cluster_id	The cluster ID that will be used to create the next cluster on this enclosure; blank if a node has been removed from a cluster.

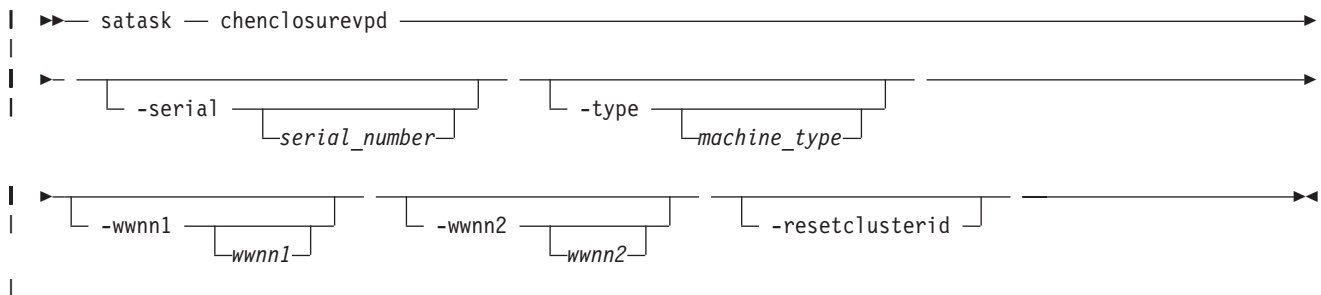
Chapter 28. Service task commands

Service information commands enable you to service the node hardware.

chenclosurevpd

Use the **chenclosurevpd** command to change fields in the control enclosure VPD.

Syntax



Parameters

-serial *serial_number*

(Optional) The new serial number for the enclosure. The *serial_number* must be set on a replacement enclosure to match the values of the enclosure being replaced.

-type *machine_type*

(Optional) The type of machine. The *machine_type* must be set on a replacement enclosure to match the values of the enclosure being replaced.

-wwnn1 *wwnn1*

(Optional) The WWNN of canister 1. The *wwnn1* must be set on a replacement enclosure to match the values of the enclosure being replaced.

Note: If you change *wwnn1* on an operating system, you might need to also change the host and fibre channel configuration settings.

-wwnn2 *wwnn2*

(Optional) The WWNN of canister 2. The *wwnn2* must be set on a replacement enclosure to match the values of the enclosure being replaced.

Note: If you change *wwnn2* on an operating system, you might need to also change the host and fibre channel configuration settings.

-resetclusterid

(Optional) Requests the stored cluster to be zeroed.

Attention: The cluster ID indicates if the enclosure, and the drives it contains, are part of a cluster. Resetting it indicates it is no longer part of a cluster, and any data required on the drives is not required. This might result in loss of access to your data.

Description

Use this command to change fields in the control enclosure VPD. The node must be in candidate or service state when you run this command, and cannot be part of a cluster.

| An example of setting the serial number

| satask chenclosurevpd -serial 123456

| chserviceip

| Use the **chserviceip** command to set the service assistant IP address for a specific node.

| Syntax

| ►► satask — chserviceip — --serviceip—*ipv4*— [—gw—*ipv4*] [—mask—*ipv4*] [—resetpassword] [panel_name] ►►

| ►► satask — chserviceip — --serviceip_6—*ipv6*— [—gw_6—*ipv6*] [—prefix_6—*int*] [—resetpassword] [panel_name] ►►

| ►► satask — chserviceip — --default— [—resetpassword] [panel_name] ►►

| ►► satask — chserviceip — --dhcp— [—resetpassword] [panel_name] ►►

| ►► satask — chserviceip — --dhcp_6— [—resetpassword] [panel_name] ►►

| Parameters

| panel_name

| (Optional) Identifies the node being serviced.

| -serviceip

| The IPv4 address for the service assistant.

| **Note:** The IPv4 service address can be unconfigured by setting the address to 0.0.0.0.

| -gw

| The IPv4 gateway for the service assistant.

| -mask

| The IPv4 subnet for the service assistant.

| -serviceip_6

| The IPv6 address for the service assistant.

| **Note:** The IPv6 service address can be unconfigured by setting the address to 0:0:0:0:0:0:0:0.

| -gw_6

| The IPv6 gateway for the service assistant.

| -prefix_6

| The IPv6 prefix for the service assistant.

| -dhcp

| Attempts to obtain an IPv4 address from DHCP.

| -dhcp_6

| Attempts to obtain an IPv6 address from DHCP.

| **-resetpassword**
| Sets the service assistant password to default.

| **Storwize V7000:**

| **panel_name**
| (Optional) Identifies the node being serviced.

| **-serviceip**
| The IPv4 address for the service assistant.

| **Note:** The IPv4 service address can be unconfigured by setting the address to 0.0.0.0.

| **-gw**
| The IPv4 gateway for the service assistant.

| **-mask**
| The IPv4 subnet for the service assistant.

| **-serviceip_6**
| The IPv6 address for the service assistant.

| **Note:** The IPv6 service address can be unconfigured by setting the address to 0:0:0:0:0:0:0:0.

| **-gw_6**
| The IPv6 gateway for the service assistant.

| **-prefix_6**
| The IPv6 prefix for the service assistant.

| **-default**
| The IPv6 prefix for the service assistant.

| **-resetpassword**
| Sets the service assistant password to default.

| **Description**

| This command sets the service assistant IP address for a specific node. If the node is part of cluster, then the cluster gateway, subnet and prefix will be used unless specified otherwise. If the node is a candidate node, the subnet, prefix and gateway must be specified. If you specify an IPV4 or IPV6 address, but do not provide a gateway, mask, or prefix, then the existing gateway, mask, and prefix values will be preserved.

| **Input**

| `satask chserviceip`

| **Input (with a specific -serviceip, -gw, and -mask parameters)**

| `satask chserviceip -serviceip 1.2.3.4 -gw 1.2.3.1 -mask 255.255.255.0`

| **chwwnn**

| Use the **chwwnn** command to modify the node World Wide Node Name (WWNN).

| **Syntax**

| `▶▶ satask — chwwnn — — -wwnnsuffix —wwnn_suffix— — [-panel_name]`

| **Parameters**

- | **-wwnnsuffix**
| (Optional) The suffix to be used for node *wwnn*.
- | **panel_name**
| (Optional) Identifies the node being serviced.

| **Description**

- | This command modifies the WWNN. Use the **lsservicestatus** command to view suggested WWNNs.
- | [No feedback]

| **cpfiles**

- | Use the **cpfiles** command to copy files from another node.

| **Syntax**

```
| ▶▶— satask — cpfiles — -prefix — directory — -source — source_panel_name —▶▶  
|                               | file_filter |  
|  
| ▶ target_panel_name ▶▶
```

| **Parameters**

- | *target_panel_name*
| Identifies the node that files will be copied to. If no panel name is provided, the files will be copied to the local node.
- | **-prefix**
| Specifies the directory, files, or directory and files to be retrieved. The path must exist in a permitted listable directory. You can use the following **-prefix** filters:
 - | • /dumps (retrieves all files in all subdirectories)
 - | • /dumps/audit
 - | • /dumps/cimom
 - | • /dumps/configs
 - | • /dumps/drive
 - | • /dumps/elogs
 - | • /dumps/enclosure
 - | • /dumps/feature
 - | • /dumps/iostats
 - | • /dumps/iotrace
 - | • /dumps/mdisk
 - | • /dumps/syslogs
 - | • /home/admin/upgrade
 - | • * /dumps/enclosure

| **Note:**

- | • You can also specify a file filter. For example, if you specify */dumps/elogs/*.txt*, all files in the */dumps/elogs* directory that end in *.txt* are copied.

- If you use a wildcard, the following rules apply:
 1. The wildcard character is an asterisk (*).
 2. The command can contain a maximum of one wildcard.
 3. When you use a wildcard, you must surround the filter entry with double quotation marks ("x"). For example: `satask cpfiles -prefix "/dumps/elogs/*.txt"`

| *source_panel_name*
 | Identifies the source node files will be copied from.

| **Description**

| This command copies files from another node. You can monitor the progress of the copy using the **sainfo** and **lscmdstatus** command.

| **To get configuration information from canister 1 in enclosure 2:**

| `satask cpfiles -prefix /dumps/configs -source 02-1`

| **installsoftware**

| Use the **installsoftware** command to install a specific software package.

| **Syntax**

| `▶▶ satask — installsoftware — — -file filename — [-ignore | -pacedccu] [panel_name]`

| **Parameters**

| **panel_name**
 | (Optional) Identifies the node being serviced.

| **-file**
 | The file name of software installation package.

| **-ignore**
 | Overrides prerequisite checking and forces installation of the software.

| **-pacedccu**
 | Causes the node to initiate a paced ccu (in which you define when the node begins its upgrade) instead of a normal ccu (in which each node in the cluster automatically upgrades in sequence).

| **Description**

| This command installs a specific software package.

| **An invocation example:**

| [No feedback]

| **leavecluster**

| Use the **leavecluster** command to remove cluster state data, location information, and other history from a node.

| **Syntax**

```
| ►► satask — leavecluster — — -force — — [ -panel_name ] —————▶▶
```

Parameters

-force

Attention: The **-force** parameter is required because this service action can cause temporary or permanent loss of access to data. Use this command only when a service procedure instructs you to do so.

panel_name

(Optional) Identifies the node being serviced.

Description

Use this command to remove cluster state data, location information, and other history from a node.

```
satask leavecluster -force
```

metadata

Use the **metadata** command to recover a virtualization table.

Syntax

```
| ►► satask — metadata — ---rebuildcluster —————▶▶
```

```
| ►► satask — metadata — ---scan— ---file—filename_arg— ---disk—UID_arg— ---start—start_arg—▶▶
```

```
| ► [ -end—end_arg ] —————▶▶
```

```
| ►► satask — metadata — ---dump— ---disk—UID_arg— ---start—start_arg —————▶▶
```

Parameters

-rebuildcluster

Creates a cluster from the metadata found in `/dumps/t3_recovery.bin` created by the `-dump` process.

-file

Specifies the file in which you want the results of a scan operation. The file is placed into the node filesystem in the directory `/dumps` and can be retrieved using `scp`. The file can be subsequently cleaned using the **cleandumps** command.

-disk

Specifies the UID of the MDisk or drive that you want to scan, or remove a dump from. .

-start

- scan: the LBA at which to start scanning
- dump: the LBA at which the metadata resides (as reported in the scan file)

-end

The last LBA in which to look for metadata on the disk.

rescuenode

Use the **rescuenode** command to start automatic recovery for a specific node.

Syntax

```
▶▶ satask — rescuenode — — -force — [ -panel_name ]
```

Parameters

panel_name

(Optional) Identifies the node being serviced.

-force

Attention: The **-force** parameter is required because this service action can cause temporary or permanent loss of access to data. Use this command only when the node reports corrupted software.

Description

This command start automatic recovery for a specific node. Use this command only when the node reports corrupted software.

```
satask rescuenode -force 112233
```

resetpassword

Use the **resetpassword** command to reset the cluster superuser password to *passw0rd*.

Syntax

```
▶▶ satask — resetpassword —
```

Parameters

None.

Description

This command resets the cluster superuser password to *passw0rd*. The next time you log in to the GUI, you will be prompted for a new password.

No feedback

snap

Use the **snap** command to create snap file on the node you specify.

Syntax

```
▶▶ satask — snap — [ -dump ] [ -panel_name ]
```


| Parameters

- | **-dump**
| (Optional) Collects the existing dump.
- | **panel_name**
| (Optional) Identifies the node being serviced.

| Description

- | This command creates a snap file on the node you specify.
- | The name of the output file is `snap.single.nodeid.date.time.tgz`.

| **setpacedccu**

- | Use the **setpacedccu** command to flag a node to participate in a user-paced concurrent code upgrade.

| Syntax

```
| ▶▶ satask — setpacedccu — panel_name ▶▶
```

| Parameters

- | **panel_name**
| (Optional) Identifies the node being serviced.

| Description

- | Use this command to flag a node to participate in a user-paced concurrent code upgrade. This command can only be used when the node is:
 - | • In a service state
 - | • Error-free
 - | • Not part of a cluster when the node is out of a service state
- | [No feedback]

| **settempsshkey**

- | Use the **settempsshkey** command to install a temporary Secure Shell (SSH) key for a superuser ID to run commands in the service assistant CLI.

| Syntax

```
| ▶▶ satask — settempsshkey — -keyfile filename — -panel_name ▶▶
```

| Parameters

- | **-keyfile *filename***
| Specifies the name of the file that contains the Secure Shell (SSH) public key. The file identified by *filename* must be on the local node (or on the USB flash drive, if you execute the command from there).

| **panel_name**
| (Optional) Identifies the node being serviced.

| **Description**

| This command installs a temporary SSH key for a superuser ID to run commands in the service assistant CLI (for example, to copy files to or from the node).


| You can only perform this command when performing a service action. Installing a temporary key will replace any available existing keys. The key will be deleted when the node joins a cluster or is rebooted or power cycled.

| [No feedback.]

| **startservice**

| Use the **startservice** command to enter a service state.

| **Syntax**

| ▶▶ satask — startservice — 

| **Parameters**

| **-force**
| (Optional) Overrides checking of cluster membership.

| **panel_name**
| (Optional) Identifies the node being serviced.

| **Description**

| This command enters a service state. For example, you might use a service state to remove a node from the candidate list, or to prevent it from automatically being added to a cluster again. The **-force** flag is required if the action could interrupt I/O (last node in cluster or IO group). This command holds the node in service state until it is cleared using the **satask stopservice** command, or until the I/O process is restarted.

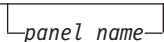
| **Input**

| satask startservice

| **stopnode**

| Use the **stopnode** command to power off, reboot, or warmstart a node.

| **Syntax**

| ▶▶ satask — stopnode — --poweroff | -reboot | -warmstart 

| **Parameters**

| **panel_name**
| (Optional) Identifies the node being serviced.

- | **-prepare**
- | Search for T3 recovery data. This locates the date of the necessary backup file and quorum disk. You can retrieve this data using the **sainfo lscmdstatus** command.
- | **-execute**
- | Start T3 recovery using recovered data.

| **Description**

- | This command prepares and starts a T3 recovery.
- | [No feedback]

|

Chapter 29. Controller command

The controller command modifies the name of a storage controller.

chcontroller

The **chcontroller** command modifies the attributes of a controller.

Syntax

```
svctask -- chcontroller -- [-name new_name] [-allowquorum yes|no]
controller_id | controller_name
```

Parameters

-name *new_name*

(Optional) Specifies the new name to be assigned to the controller.

-allowquorum *yes* | *no*

(Optional) Specifies that the controller is allowed or is not allowed to support quorum disks. A value of **yes** enables a suitable controller to support quorum disks. A value of **no** disables a controller from supporting quorum disks, provided that the specified controller is not currently hosting a quorum disk.

controller_id | *controller_name*

(Required) Specifies the controller to modify; use either the controller name or the controller ID.

Description

This command changes the name of the controller that is specified by the *controller_id* | *controller_name* variable to the value that you specify with the **-name** parameter.

If any controller that is associated with an MDisk shows the **allow_quorum** attribute set to no with the **svcinfolcontroller** command, the set quorum action fails for that MDisk. Before using the **svctask chcontroller** command to set the **-allowquorum** parameter to yes on any disk controller, check the following Web site to see whether the controller supports quorum.

Support for SAN Volume Controller (2145) website at www.ibm.com/storage/support/2145

You can add a new disk controller system to your SAN at any time. Follow the switch zoning guidelines in the section about switch zoning. Also, ensure that the controller is set up correctly for use with the cluster.

To add a new disk controller system to a running configuration, ensure that the cluster has detected the new storage MDisks by issuing the **svctask detectmdisk** command. The controller has automatically been assigned a default name. If you are unsure of which controller is presenting the MDisks, issue the **svcinfolcontroller** command to list the controllers. The new controller is listed with the highest numbered default name. Record the controller name and follow the instructions in the section about determining a disk controller system name.

- 1 Give this controller a descriptive name by issuing the following command:

```
svctask chcontroller -name newname oldname
```

List the unmanaged MDisks by issuing the following command:

```
svcinfolsmdisk -filtervalue mode=unmanaged:controller_name=newname
```

These MDisks correspond to the RAID arrays or partitions that you have created. Record the field controller LUN number. The field controller LUN number corresponds with the LUN number that you assigned to each of the arrays or partitions.

Create a new managed disk group and add only the RAID arrays that belong to the new controller to this MDisk group. Avoid mixing RAID types; for each set of RAID array types (for example, RAID-5 or RAID-1), create a new MDisk group. Assign this MDisk group an appropriate name; if your controller is called FAST650-abc and the MDisk group contains RAID-5 arrays, assign the MDisk a name similar to **F600-abc-R5**. Issue the following command:

```
svctask mkmdiskgrp -ext 16 -name mdisk_grp_name  
-mdisk colon-separated list of RAID-x mdisks returned
```

Note: This creates a new MDisk group with an extent size of 16 MB.

An invocation example

```
svctask chcontroller -name newtwo 2
```

The resulting output

No feedback

CMMVC5000I No message was found for major rc MAJOR_RC , minor rc MINOR_RC , for action/view id ACTION_VIEW_ID .

Explanation: A message is missing.

User response: Contact the support center.

CMMVC5700E The parameter list is not valid.

Explanation: You have entered a list of parameters that is not supported for the command.

User response: Specify a parameter list that is supported for the command, and resubmit the command.

CMMVC5701E No object ID was specified.

Explanation: The command that you have submitted requires that you specify an object identifier name or ID number, and you did not specify an object identifier.

User response: Specify an object ID, and resubmit the command.

CMMVC5702E VALUE is below the minimum level.

Explanation: You entered the specified string as a value for a parameter. The parameter requires a

minimum value, and the specified string is less than the required minimum value.

User response: Specify a value that is supported by the parameter, and resubmit the command.

CMMVC5703E The value or list starting with VALUE is above the maximum permitted for that value or has exceeded the number of items allowed in a list.

Explanation: You have entered the specified string as a value for a parameter. The string is either a standalone value or the first value in a list of values. If the string is a standalone value, the value is greater than the supported maximum value for the parameter. If the string is the first value in a list of values, the list contains more than the supported maximum number of entries for the parameter.

User response: Specify a value or list of values that is supported by the parameter, and resubmit the command.

CMMVC5704E VALUE is not divisible by the permitted step value.

Explanation: You have entered the specified string as a value for a parameter. The string is not a supported value for the parameter. One requirement is that the

value is an even multiple of 16, and the specified string does not meet that requirement.

User response: Specify a value that is supported by the parameter, and resubmit the command.

CMMVC5705E A required parameter is missing.

Explanation: The command that you have submitted has at least one required parameter that you have not entered.

User response: Specify all of the required parameters, and resubmit the command.

CMMVC5706E An invalid argument has been entered for the *PARAMETER* parameter.

Explanation: You have entered a value for the specified parameter and the value is not supported for the parameter. The parameter supports a specific set of values.

User response: Specify a value that is supported by the parameter, and resubmit the command.

CMMVC5707E Required parameters are missing.

Explanation: The command that you have submitted has more than one required parameter that you have not entered.

User response: Specify all of the required parameters, and resubmit the command.

CMMVC5708E The *PARAMETER* parameter is missing its associated arguments.

Explanation: You have entered the specified parameter without an associated value. This parameter, like most parameters, requires an associated value.

User response: Specify the associated value, and resubmit the command.

CMMVC5709E *VALUE* is not a supported parameter.

Explanation: The specified string is not a supported parameter for the command that you have entered.

User response: Specify the correct parameter, and resubmit the command.

CMMVC5711E *VALUE* is not valid data.

Explanation: You have entered the specified string as a value for a parameter. The string is not a supported value for the parameter.

User response: Specify a value that is supported by the parameter, and resubmit the command.

CMMVC5712E Required data is missing.

Explanation: You have entered an incomplete command.

User response: Specify command completely, and resubmit the command.

CMMVC5713E Some parameters are mutually exclusive.

Explanation: Certain commands have two or more parameters that are mutually exclusive. You have submitted a command using at least two mutually exclusive parameters.

User response: Specify a supported combination of parameters, and resubmit the command.

CMMVC5714E The parameter list is empty.

Explanation: Certain parameters require one or more values in a colon separated parameter list. You have specified at least one parameter without the required parameter list.

User response: Specify at least one value for all parameters that require a value, and resubmit the command.

CMMVC5715E The parameter list does not exist.

Explanation: Certain parameters require one or more values in a colon separated parameter list. You have specified at least one parameter without the required parameter list.

User response: Specify at least one value for all parameters that require a value, and resubmit the command.

CMMVC5716E Non-numeric data was entered for the numeric field *FIELD*. Enter a numeric value.

Explanation: You have entered the specified string as a value for a parameter that supports only numeric values.

User response: Specify a numeric value in the numeric field, and resubmit the command.

CMMVC5717E No match was found for the specified unit.

Explanation: Certain parameters allow a user to specify a data unit such as mb or kb. You have entered a data unit for a parameter that supports data units, but the data unit that you have entered is not a supported data unit for the parameter.

User response: Specify the correct data unit, and resubmit the command.

CMMVC5718E An unexpected return code was received.

Explanation: The command has completed, but the acknowledgement of the command completion contains a return code that is not defined.

User response: Determine whether or not the command has succeeded. If the command has not succeeded, resubmit the command. If the problem persists, contact IBM technical support for assistance.

CMMVC5719E A value of *VALUE* requires the parameter *PARAMETER* to be specified.

Explanation: Certain commands have required combinations of parameters based on either the entry of a parameter or the value for a parameter. When you enter the specified value, you must enter the specified parameter.

User response: Specify the required parameter, and resubmit the command.

CMMVC5721E *VALUE* is not a valid time stamp format. The valid format is MMDDHHmmYYYY.

Explanation: The specified value is not a valid time-stamp format. The valid format is MMDDHHmmYYYY.

User response: Follow the correct time-stamp format, and resubmit the command.

CMMVC5722E *VALUE* is not a valid month.

Explanation: The specified value is not a valid month.

User response: Specify the correct month (MM), and resubmit the command.

CMMVC5723E *VALUE* is not a valid day.

Explanation: The specified value is not a valid day.

User response: Specify the correct day (DD), and resubmit the command.

CMMVC5724E *VALUE* is not a valid hour.

Explanation: The specified value is not a valid hour.

User response: Specify the correct hour (HH), and resubmit the command.

CMMVC5725E *VALUE* is not a valid minute.

Explanation: The specified value is not a valid minute.

User response: Specify the correct minute (mm), and resubmit the command.

CMMVC5726E *VALUE* are not valid seconds.

Explanation: The specified value are not valid seconds.

User response: Specify the correct seconds (ss), and resubmit the command.

CMMVC5727E *VALUE* is not a valid filter.

Explanation: You can filter the output of some views by using the *-filtervalue* parameter. The specified string that you have entered is not a supported value for the *-filtervalue* parameter in this view.

User response: Ensure that you use a supported value for the *-filtervalue* parameter, and resubmit the command.

CMMVC5728E *VALUE* should be in the format minute:hour:day:month:weekday.

Explanation: The specified value should be in the format minute:hour:day:month:weekday.

User response: Follow the correct format, and resubmit the command.

CMMVC5729E One or more components in the list is not valid.

Explanation: Certain parameters support one or more items of data in a colon separated list. At least one of the items in the list that you have entered is not correct.

User response: Ensure that you enter supported values in the list, and resubmit the command.

CMMVC5730E *VALUE* is only valid when *VALUE* has a value of *VALUE* .

Explanation: The specified command and parameter combination that you have entered requires the specified parameter value.

User response: Ensure that you specify the correct parameter value for the command and parameter combination that you enter, and resubmit the command.

CMMVC5731E *VALUE* can only be entered when *VALUE* has been entered.

Explanation: Certain commands have required combinations of parameters based either on the inclusion of a specified parameter, or on the value entered for a specified parameter. When you include the first specified string in the command, you must enter the second specified string as a parameter.

User response: Ensure that you enter a supported

combination of parameters and values, and resubmit the command.

CMMVC5732E The command cannot be initiated because it was not run on the configuration node.

Explanation: The `svctask` and `svcinfo` commands must be run on the configuration node.

User response: Log off of the node service IP address, log on to the cluster IP address, and run the command on the configuration node.

CMMVC5733E Enter at least one parameter.

Explanation: You must specify at least one parameter for the command that you have submitted.

User response: Specify at least one parameter, and resubmit the command.

CMMVC5734E A combination of values was entered that is not valid.

Explanation: You have specified a combination of values that is not correct.

User response: Specify a supported combination of values, and resubmit the command.

CMMVC5735E The name entered is not valid. Enter an alphanumeric string that does not start with a number.

Explanation: The first character of an object name cannot be numeric.

User response: Specify an alphanumeric string that does not start with a numeric, and resubmit the command.

CMMVC5737E The parameter *PARAMETER* has been entered multiple times. Enter the parameter only one time.

Explanation: The specified parameter was entered more than once.

User response: Delete all duplicate parameters, and resubmit the command.

CMMVC5738E The argument *ARGUMENT* contains too many characters.

Explanation: The field length of the specified argument is longer than the maximum supported field length for the argument.

User response: Specify the correct argument, and resubmit the command.

CMMVC5739E The argument *ARGUMENT* does not contain enough characters.

Explanation: The field length of the specified argument is less than the minimum supported field length for the argument.

User response: Specify the correct argument, and resubmit the command.

CMMVC5740E The filter flag *VALUE* is not valid.

Explanation: You can filter the output of some views by using the `-filtervalue` parameter. The specified string that you have entered is not a supported value for the `-filtervalue` parameter in this view.

User response: Ensure that you use a supported value for the `-filtervalue` parameter, and resubmit the command.

CMMVC5741E The filter value *VALUE* is not valid.

Explanation: You can filter the output of some views by using the `-filtervalue` parameter. Each filter has an associated value. The syntax is `-filtervalue filter=value`. The specified string that you have entered is not a supported value for the `-filtervalue filter` that you specified in this view.

User response: Ensure that you use a supported value for the `-filtervalue filter` that you specify, and resubmit the command.

CMMVC5742E A specified parameter is out of its valid range.

Explanation: You have entered data that is not in the range of values that is supported for the parameter that you have entered.

User response: Ensure that you enter data values that are supported for the parameter that you enter, and resubmit the command.

CMMVC5743E A specified parameter does not comply with the step value.

Explanation: A parameter was specified that does not comply with the step value.

User response: Specify the correct parameter, and resubmit the command.

CMMVC5744E Too many objects were specified in the command.

Explanation: There were too many objects specified in the command.

User response: Specify the correct object, and resubmit the command.

CMMVC5745E Too few objects were specified in the request.

Explanation: There were not enough objects specified in the command.

User response: Specify the correct object, and resubmit the command.

CMMVC5746E The requested operation cannot be applied to the object specified.

Explanation: The requested operation is not valid for this object.

User response: Specify a valid operation, and resubmit the command.

CMMVC5747E The action requested is invalid - internal error.

Explanation: The operation that was requested is not valid.

User response: Specify the correct operation, and resubmit the command.

CMMVC5748E The action requested is invalid - internal error.

Explanation: The operation that was requested is not valid.

User response: Specify the correct operation, and resubmit the command.

CMMVC5749E The dump filename specified already exists.

Explanation: The dump file name that was specified already exists.

User response: Specify a different dump file name, and resubmit the command.

CMMVC5750E The dump file could not be created - the file system is probably full.

Explanation: The dump file was not created. The file system might be full.

User response: Not applicable.

CMMVC5751E The dump file could not be written to.

Explanation: The dump file could not be written to disk.

User response: Not applicable.

CMMVC5752E Request failed. The object contains child objects, these must be deleted first.

Explanation: The operation failed because the specified object contains child objects.

User response: Delete the child objects, and resubmit the command.

CMMVC5753E The specified object does not exist or is not a suitable candidate.

Explanation: The specified object does not exist or is not a suitable candidate.

User response: Specify the correct object, and resubmit the command.

CMMVC5754E The specified object does not exist, or the name supplied does not meet the naming rules.

Explanation: The specified object does not exist, or the name of the object does not meet the naming requirements.

User response: Specify the correct object name, and resubmit the command.

CMMVC5755E Cannot create as the sizes of the specified objects do not match.

Explanation: The sizes of the specified objects do not match.

User response: Not applicable.

CMMVC5756E Cannot perform the request as the object id is already mapped to another object or is the subject of an FC or RC relationship.

Explanation: The operation failed because the specified object is already mapped.

User response: Specify a different object, and resubmit the command.

CMMVC5757E Self Defining Structure (SDS) defaults not found - internal error.

Explanation: The defaults for the self describing structure were not found.

User response: Not applicable.

CMMVC5758E Object name already exists.

Explanation: The object name already exists.

User response: Specify a unique object name, and resubmit the command.

CMMVC5759E An internal error has occurred - memory could not be allocated.

Explanation: The memory cannot be allocated.

User response: Not applicable.

CMMVC5760E Failed to add the node to the cluster member list.

Explanation: The node could not be added to the cluster.

User response: Not applicable.

CMMVC5761E Failed to delete the node from the cluster member list.

Explanation: The node could not be deleted from the cluster.

User response: Not applicable.

CMMVC5762E The request did not complete before the timeout period expired.

Explanation: The operation failed because the timeout period expired.

User response: Resubmit the command.

CMMVC5763E The node failed to go online.

Explanation: The node failed to go online.

User response: Not applicable.

CMMVC5764E The mode change request is invalid - internal error

Explanation: The specified mode change is not valid.

User response: Specify a different mode, and resubmit the command.

CMMVC5765E The object specified is no longer a candidate - a change occurred during the request.

Explanation: The specified object is no longer a candidate. A change occurred during the request.

User response: Specify a different object, and resubmit the command.

CMMVC5767E One or more of the parameters specified are invalid or a parameter is missing.

Explanation: One or more of the specified parameters is not valid.

User response: Specify the correct parameter, and resubmit the command.

CMMVC5769E The requested operation requires all nodes to be online - one or more nodes are not online.

Explanation: The operation requires that all nodes be online. One or more nodes are not online.

User response: Check that each node is online, and resubmit the command.

CMMVC5770E The SSH key file supplied is invalid.

Explanation: The file for the SSH key is not valid.

User response: Specify a different file, and resubmit the command.

CMMVC5771E The operation requested could not complete, usually due to child objects existing. To force the operation, specify the force flag.

Explanation: The operation failed, probably, because the object contains child objects.

User response: Specify the -force flag to complete the operation, and resubmit the command.

CMMVC5772E The operation requested could not be performed because software upgrade is in progress.

Explanation: The operation failed because a software upgrade is in progress.

User response: Wait for the software upgrade to complete, and resubmit the command.

CMMVC5773E The object selected is in the wrong mode to perform the requested operation.

Explanation: The operation failed because the selected object is in the wrong mode.

User response: Specify the correct mode, and resubmit the command.

CMMVC5774E The userid supplied is not valid.

Explanation: The userid is not valid.

User response: Specify a different userid, and resubmit the command.

CMMVC5775E The directory attribute specified is not valid.

Explanation: The directory attribute is not valid.

User response: Specify a different directory, and resubmit the command.

CMMVC5776E The directory listing could not be retrieved.

Explanation: The directory listing could not be retrieved.

User response: Specify a different directory listing, and resubmit the command.

CMMVC5777E The node could not be added to the IO Group, because the other node in the IO Group is in the same power domain.

Explanation: The node was not added to the I/O group because the other node in the I/O Group is in the same power domain.

User response: Specify a different node from another I/O group, and resubmit the command.

CMMVC5778E Cannot create another cluster, a cluster already exists.

Explanation: The cluster was not created because one already exists.

User response: Not applicable.

CMMVC5780E The action could not be completed using the Remote Cluster name. Use the Remote Cluster Unique ID instead.

Explanation: The unique ID of the remote cluster is required for this command.

User response: Specify the unique ID of the remote cluster, and resubmit the command.

CMMVC5781E The cluster ID specified is invalid.

Explanation: The cluster ID is not valid.

User response: Specify a different cluster ID, and resubmit the command.

CMMVC5782E The object specified is offline.

Explanation: The object is offline.

User response: Specify an object that is online, and resubmit the command.

CMMVC5783E The information is not available to complete this command.

Explanation: This error is only returned when the node is in the service state.

User response: None.

CMMVC5784E The cluster name specified is not unique, specify the cluster using the cluster ID.

Explanation: The cluster name is not unique.

User response: Specify the cluster using the cluster ID, and resubmit the command.

CMMVC5785E The filename specified contains an illegal character.

Explanation: The filename contains an illegal character.

User response: Specify a valid filename, and resubmit the command.

CMMVC5786E The action failed because the cluster is not in a stable state.

Explanation: The action failed because the cluster is not in a stable state.

User response: Not applicable.

CMMVC5787E The cluster was not created because a cluster already exists.

Explanation: The cluster was not created because a cluster already exists.

User response: Not applicable.

CMMVC5788E The service IP address is not valid.

Explanation: The service IP address is not valid.

User response: Specify the correct service IP address, and resubmit the command.

CMMVC5789E The cluster was not modified because the IP address, subnet mask, service address, SNMP address, or gateway address is not valid.

Explanation: The cluster was not modified because the IP address, subnet mask, service address, SNMP address, or gateway address is not valid.

User response: Specify all correct attributes, and resubmit the command.

CMMVC5790E The node was not added to the cluster because the maximum number of nodes has been reached.

Explanation: The node was not added to the cluster because the maximum number of nodes has been reached.

User response: Not applicable.

CMMVC5791E The action failed because an object that was specified in the command does not exist.

Explanation: An entity that was specified in the command does not exist, therefore the action failed.

User response: Specify the correct entity, and resubmit the command.

CMMVC5792E The action failed because the I/O group is used for recovery.

Explanation: The action failed because the I/O group is used for recovery.

User response: Not applicable.

CMMVC5793E The node was not added to the cluster because the I/O group already contains a pair of nodes.

Explanation: The node was not added to the cluster because the I/O group already contains a pair of nodes.

User response: Not applicable.

CMMVC5794E The action failed because the node is not a member of the cluster.

Explanation: The node is not a member of the cluster, therefore the action failed.

User response: Specify a node that is contained in the cluster, and resubmit the command.

CMMVC5795E The node was not deleted because a software upgrade is in progress.

Explanation: The node was not deleted because a software upgrade is in progress.

User response: Wait for the software upgrade to complete, and resubmit the command.

CMMVC5796E The action failed because the I/O group that the node belongs to is unstable.

Explanation: A previous configuration command might not yet have completed.

User response: Wait for the previous command to complete, and resubmit the command.

CMMVC5797E The node was not deleted because this is the last node in the I/O group and there are virtual disks (VDisks) associated with the I/O group.

Explanation: The specified node is the last node in the I/O group and there are volumes associated with the

I/O group, therefore the node could not be deleted.

User response: Not applicable.

CMMVC5798E The action failed because the node is offline.

Explanation: The action failed because the node is offline.

User response: Specify a node that is online, and resubmit the command.

CMMVC5799E The shut down was not successful because there is only one online node in the I/O group.

Explanation: There is only one online node in the I/O group, therefore the shut down operation was not successful.

User response: Not applicable.

CMMVC5800E The action failed because an entity that was specified in the command does not exist.

Explanation: The entity that was specified in the command does not exist, therefore the action failed.

User response: Specify a different entity, and resubmit the command.

CMMVC5801E The upgrade of the cluster software could not proceed because every node in the cluster must be online. Either delete the node that is offline or bring the node online and resubmit the command.

Explanation: The upgrade of the cluster software could not proceed because every node in the cluster must be online.

User response: Either delete the node that is offline or bring the node online, and resubmit the command.

CMMVC5802E The upgrade of the cluster software could not proceed because there is an I/O group in the cluster that contains only one node. The software upgrade requires that each node in an I/O group be shut down and restarted. If there is only one node in an I/O group, I/O operations could be lost if I/O operations are not stopped before beginning the software upgrade.

Explanation: The upgrade of the cluster software could not proceed because there is an I/O group in the cluster that contains only one node. The software upgrade requires that each node in an I/O group be shut down and restarted. If there is only one node in an I/O group, I/O operations could be lost if I/O

CMMVC5803E • CMMVC5814E

operations are not stopped before beginning the software upgrade.

User response: Either upgrade the cluster using the -force option or specify a different node, and resubmit the command.

CMMVC5803E The entry in the error log was not marked because the error is already fixed or unfixed, or the sequence number could not be found.

Explanation: The entry in the event log was not marked because the sequence number was not found.

User response: Not applicable.

CMMVC5804E The action failed because an object that was specified in the command does not exist.

Explanation: The entity that was specified in the command does not exist, therefore the action failed.

User response: Specify a different entity, and resubmit the command.

CMMVC5805E The progress information was not returned because the FlashCopy statistics are not ready yet.

Explanation: The progress information was not returned because the FlashCopy statistics are not ready yet.

User response: Not applicable.

CMMVC5806E The action failed because an object that was specified in the command does not exist.

Explanation: The entity that was specified in the command does not exist, therefore the action failed.

User response: Specify a different entity, and resubmit the command.

CMMVC5807E The action failed because the managed disk (MDisk) cannot be changed to the specified mode.

Explanation: The action failed because the managed disk (MDisk) cannot be changed to the specified mode.

User response: Not applicable.

CMMVC5808E The action failed because the managed disk (MDisk) does not exist.

Explanation: The action failed because the managed disk (MDisk) does not exist.

User response: Specify a different MDisk, and resubmit the command.

CMMVC5809E The tracing of I/O operations was not started because it is already in progress.

Explanation: The tracing of I/O operations was not started because it is already in progress.

User response: Not applicable.

CMMVC5810E The quorum index number for the object was not set because the object is not available, either because the object is offline or because it is associated with another quorum disk.

Explanation: The object that you specify for this task must be online and cannot already be a quorum disk.

User response: Change the status of the specified object, or specify a different object, and resubmit the command.

CMMVC5811E The quorum index number for the object was not set because the quorum disk does not exist.

Explanation: An existing quorum disk must be specified before the quorum index number of the object can be set.

User response: Specify an existing quorum disk, and resubmit the command.

CMMVC5812E The quorum index number for the object was not set because the object is in the wrong mode.

Explanation: The object must be a managed MDisk or a candidate disk before the quorum index number of the object can be set.

User response: Change the mode of the specified object, or specify a different object, and resubmit the command.

CMMVC5813E The quorum index number for the object was not set because the object has a sector size that is not valid.

Explanation: The sector size of the specified object will not allow the quorum index number for the object to be set.

User response: Change the sector size of the specified object, or specify a different object, and resubmit the command.

CMMVC5814E The quorum index number for the managed disk (MDisk) was not set because quorum is not allowed on one or more associated controllers.

Explanation: The quorum index number for the managed disk (MDisk) was not set because quorum is

not allowed on one or more associated controllers.

User response: Specify an MDisk that has quorum enabled on all of its associated controllers, and resubmit the command.

CMMVC5815E The managed disk group was not created because an entity that was specified in the command does not exist.

Explanation: The storage pool was not created because an entity that was specified in the command does not exist.

User response: Specify a different entity, and resubmit the command.

CMMVC5816E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

User response: Specify a different entity, and resubmit the command.

CMMVC5817E The specified managed disk group was invalid.

Explanation: The storage pool was not renamed because the name was not valid.

User response: Specify a different storage pool name, and resubmit the command.

CMMVC5818E The managed disk group was not deleted because there is at least one MDisk in the group.

Explanation: The storage pool was not deleted because there is at least one MDisk in the group.

User response: Not applicable.

CMMVC5819E The managed disk (MDisk) was not added to the MDisk group because the MDisk is part of another MDisk group.

Explanation: The managed disk (MDisk) was not added to the storage pool because the MDisk is part of another storage pool.

User response: Not applicable.

CMMVC5820E The managed disk (MDisk) was not added to the MDisk group because an entity that was specified in the command does not exist.

Explanation: The managed disk (MDisk) was not added to the storage pool because an entity that was

specified in the command does not exist.

User response: Specify a different entity, and resubmit the command.

CMMVC5821E The managed disk (MDisk) was not added to the MDisk group because not enough MDisks were included in the list.

Explanation: The managed disk (MDisk) was not added to the storage pool because not enough MDisks were included in the list.

User response: Include more MDisks in the list, and resubmit the command.

CMMVC5822E The managed disk (MDisk) was not added to the MDisk group because too many MDisks were included in the list.

Explanation: The managed disk (MDisk) was not added to the storage pool because too many MDisks were included in the list.

User response: Delete the extra MDisks in the list, and resubmit the command.

CMMVC5823E The managed disk (MDisk) was not deleted from the MDisk group because the MDisk is part of another MDisk group.

Explanation: The managed disk (MDisk) was not deleted from the storage pool because the MDisk is part of another storage pool.

User response: Not applicable.

CMMVC5824E The managed disk (MDisk) was not deleted from the MDisk group because it does not belong to the MDisk group.

Explanation: The managed disk (MDisk) was not deleted from the storage pool because it does not belong to the storage pool.

User response: Not applicable.

CMMVC5825E The managed disk (MDisk) was not deleted from the MDisk group because a virtual disk (VDisk) is allocated from one or more of the specified MDisks. A forced deletion is required.

Explanation: The managed disk (MDisk) was not deleted from the storage pool because a volume is allocated from one or more of the specified MDisks.

User response: Specify the -force option, and resubmit the command.

CMMVC5826E The virtual disk (VDisk) was not created because an entity that was specified in the command does not exist.

Explanation: The volume was not created because an entity that was specified in the command does not exist.

User response: Specify a different entity, and resubmit the command.

CMMVC5827E The command failed as a result of either an inconsistency between two or more of the entered parameters, or an inconsistency between a parameter and the requested action.

Explanation: The command failed as a result of an inconsistency between two or more of the entered parameters.

User response: Specify one parameter, and resubmit the command.

CMMVC5828E The virtual disk (VDisk) was not created because the I/O group contains no nodes.

Explanation: The volume was not created because the I/O group contains no nodes.

User response: Not applicable.

CMMVC5829E The image-mode or sequential-mode virtual disk (VDisk) was not created because more than one managed disk (MDisk) is specified.

Explanation: The image-mode or sequential-mode volume was not created because more than one MDisk is specified.

User response: Specify a different MDisk, and resubmit the command.

CMMVC5830E The image-mode virtual disk (VDisk) was not created because no managed disk (MDisk) was specified in the command.

Explanation: The image-mode volume was not created because no managed disk (MDisk) was specified in the command.

User response: Specify a MDisk, and resubmit the command.

CMMVC5831E The virtual disk (VDisk) was not created because the preferred node for I/O operations is not part of the I/O group.

Explanation: The volume was not created because the preferred node for I/O operations is not part of the I/O group.

User response: Specify a different node, and resubmit the command.

CMMVC5832E The property of the virtual disk (VDisk) was not modified because an entity that was specified in the command does not exist.

Explanation: The property of the volume was not modified because an entity that was specified in the command does not exist.

User response: Specify a different entity, and resubmit the command.

CMMVC5833E The property of the virtual disk (VDisk) was not modified because there are no nodes in the I/O group.

Explanation: The property of the volume was not modified because there are no nodes in the I/O group.

User response: Not applicable.

CMMVC5834E The I/O group for the virtual disk (VDisk) was not modified because the group is a recovery I/O group. To modify the I/O group, use the force option.

Explanation: The I/O group for the volume was not modified because the group is a recovery I/O group.

User response: Specify the -force option, and resubmit the command.

CMMVC5835E The virtual disk (VDisk) was not expanded because an entity that was specified in the command does not exist.

Explanation: The volume was not expanded because an entity that was specified in the command does not exist.

User response: Specify a different entity, and resubmit the command.

CMMVC5836E The virtual disk (VDisk) was not shrunk because it is locked.

Explanation: Commands might still be running in the background.

User response: Wait for all commands to complete. Use the `svcinfo lsmigrate` command to view any migrates running in the background.

CMMVC5837E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.

Explanation: The action failed because the volume is part of a FlashCopy mapping.

User response: Specify a different volume that is not part of a FlashCopy mapping, and resubmit the command.

CMMVC5838E The action failed because the virtual disk (VDisk) is part of a Remote Copy mapping.

Explanation: The action failed because the volume is part of a Remote Copy mapping.

User response: Specify a different volume that is not part of a Remote Copy mapping, and resubmit the command.

CMMVC5839E The virtual disk (VDisk) was not shrunk because an object that was specified in the command does not exist.

Explanation: The volume was not shrunk because an object that was specified in the command does not exist.

User response: Specify a different object, and resubmit the command.

CMMVC5840E The virtual disk (VDisk) was not deleted because it is mapped to a host or because it is part of a FlashCopy or Remote Copy mapping, or is involved in an image mode migrate.

Explanation: The volume was not deleted because it is mapped to a host or because it is part of a FlashCopy or Metro Mirror mapping.

User response: Specify a different volume, and resubmit the command.

CMMVC5841E The virtual disk (VDisk) was not deleted because it does not exist.

Explanation: The volume was not deleted because it does not exist.

User response: Specify a different volume, and resubmit the command.

CMMVC5842E The action failed because an object that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

User response: Specify a different entity, and resubmit the command.

CMMVC5843E The VDisk-to-host mapping was not created because the VDisk does not have a capacity greater than zero bytes.

Explanation: The host map was not created because the volume does not have a capacity greater than zero bytes.

User response: Specify a volume in which its capacity is greater than zero bytes, and resubmit the command.

CMMVC5844E The VDisk-to-host mapping was not created because the SCSI logical unit number (LUN) ID is not valid.

Explanation: The host map was not created because the SCSI logical unit number (LUN) ID is not valid.

User response: Specify the correct SCSI logical unit number (LUN) ID, and resubmit the command.

CMMVC5845E The extent was not migrated because an object that was specified in the command does not exist.

Explanation: The extent was not migrated because an object that was specified in the command does not exist.

User response: Specify a different object, and resubmit the command.

CMMVC5846E The virtual disk (VDisk) was not migrated because an object that was specified in the command does not exist.

Explanation: The volume was not migrated because an object that was specified in the command does not exist.

User response: Specify a different object, and resubmit the command.

CMMVC5847E The virtual disk (VDisk) was not migrated because its associated managed disk (MDisk) is already in the MDisk group.

Explanation: The volume was not migrated because its associated managed disk (MDisk) is already in the storage pool.

User response: Not applicable.

CMMVC5848E The action failed because the virtual disk (VDisk) does not exist or it is being deleted.

Explanation: The action failed because the volume does not exist or it is being deleted.

User response: Specify a different volume, and resubmit the command.

CMMVC5849E The migration failed because some or all of the extents are already being migrated.

Explanation: The migration failed because some or all of the extents are already being migrated.

User response: Not applicable.

CMMVC5850E The extent was not migrated because there is a problem with the source extents.

Explanation: The extent was not migrated because there is a problem with the source extents.

User response: Not applicable.

CMMVC5851E The extent was not migrated because there is a problem with the target extents.

Explanation: The extent was not migrated because there is a problem with the target extents.

User response: Not applicable.

CMMVC5852E The migration failed because there are too many migrations in progress.

Explanation: The migration failed because there are too many migrations in progress.

User response: Wait for the migration process to complete, and resubmit the command.

CMMVC5853E The action failed because there was a problem with the group.

Explanation: An attempt was made to work on a volume which is using a storage pool with one of the following problems:

- The target and source storage pools have different extent sizes (group migrate).
- The target and source storage pools are the same (group migrate).
- The target and source storage pools are different (extents migrate).
- The target group (group migrate) is not valid.

- The source group (group migrate) is not valid.

User response: Ensure that none of the above conditions exist before reissuing the command.

CMMVC5854E The extent information was not returned because the extent is not used or does not exist.

Explanation: The extent information was not returned because the extent is not used or does not exist.

User response: Specify the correct extent, and resubmit the command.

CMMVC5855E The extent information was not returned because the managed disk (MDisk) is not used by any virtual disk (VDisk).

Explanation: The extent information was not returned because the managed disk (MDisk) is not used by any volume.

User response: Specify the correct MDisk, and resubmit the command.

CMMVC5856E The action failed because the virtual disk (VDisk) does not belong to the specified managed disk group.

Explanation: The action failed because the volume does not belong to the specified storage pool.

User response: Specify a different volume, and resubmit the command.

CMMVC5857E The action failed because the managed disk (MDisk) does not exist or it is not a member of the managed disk group.

Explanation: The action failed because the managed disk (MDisk) does not exist or it is not a member of the storage pool.

User response: Specify a different MDisk, and resubmit the command.

CMMVC5858E The action failed because the virtual disk (VDisk) is in the wrong mode, the managed disk (MDisk) is in the wrong mode, or both are in the wrong mode.

Explanation: The action failed because the volume is in the wrong mode, the managed disk (MDisk) is in the wrong mode, or both are in the wrong mode.

User response: Check that the volume and MDisk are in the correct mode, and resubmit the command.

CMMVC5859E The migration did not complete because an error occurred during the migration of the last extent on an image-mode virtual disk (VDisk).

Explanation: The migration did not complete because an error occurred during the migration of the last extent on an image-mode volume.

User response: Not applicable.

CMMVC5860E The action failed because there were not enough extents in the managed disk group.

Explanation: This error is also returned if a stripe set of MDisks has been specified and one or more of these MDisks does not contain enough free extents to complete the creation of the volume.

User response: In this case, the storage pool reports that it has enough free capacity to create the volume. You can check the free capacity on each MDisk by submitting the command `svcinfolsfreeextents <mdiskname/ID>`. Alternatively, do not specify a stripe set and let the system choose the free extents automatically.

CMMVC5861E The action failed because there were not enough extents on the managed disk (MDisk).

Explanation: The action failed because there were not enough extents on the managed disk (MDisk).

User response: Specify another extent, and resubmit the command.

CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted.

Explanation: The action failed because the volume is being formatted.

User response: Wait for the volume to be successfully formatted, and resubmit the command.

CMMVC5863E The migration failed because there are not enough free extents on the target managed disk (MDisk).

Explanation: The migration failed because there are not enough free extents on the target managed disk (MDisk).

User response: Specify another free extent, and resubmit the command.

CMMVC5864E The extent information was not returned because the source extent is not used.

Explanation: The extent information was not returned because the source extent is not used.

User response: Specify a different source extent, and resubmit the command.

CMMVC5865E The action failed because the extent is out of range for the managed disk (MDisk) or virtual disk (VDisk) specified.

Explanation: The extent information was not returned because the extent is out of range for the managed disk (MDisk) or volume.

User response: Specify a different extent which is in range for the MDisk or volume and resubmit the command.

CMMVC5866E The action failed because the extent contains internal data.

Explanation: The extent was not migrated because the extent contains internal data.

User response: Not applicable.

CMMVC5867E The action failed because the worldwide port name is already assigned or is not valid.

Explanation: The action failed because the worldwide port name is already assigned or is not valid.

User response: Specify a different worldwide port name, and resubmit the command.

CMMVC5868E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

User response: Specify a different entity, and resubmit the command.

CMMVC5869E The host object was not renamed because the host ID or name is not valid.

Explanation: The host object was not renamed because the host ID or name is not valid.

User response: Specify a different host ID or name, and resubmit the command.

CMMVC5870E The host object was not deleted because an entity that was specified in the command does not exist.

Explanation: The host object was not deleted because an entity that was specified in the command does not exist.

User response: Specify the correct entity, and resubmit the command.

CMMVC5871E The action failed because one or more of the configured port names is in a mapping.

Explanation: The action failed because one or more of the configured port names is in a mapping.

User response: Specify a port name that is not in a mapping, and resubmit the command.

CMMVC5872E The port (WWPN) was not added to the host object because an object that was specified in the command does not exist.

Explanation: The port (WWPN) was not added to the host object because an object that was specified in the command does not exist.

User response: Specify the correct object, and resubmit the command.

CMMVC5873E No matching WWPN.

Explanation: The action failed because there is no matching worldwide port name.

User response: Not applicable.

CMMVC5874E The action failed because the host does not exist.

Explanation: The action failed because the host does not exist.

User response: Specify a different host, and resubmit the command.

CMMVC5875E The action failed because the virtual disk (VDisk) does not exist.

Explanation: The action failed because the volume does not exist.

User response: Specify a different volume, and resubmit the command.

CMMVC5876E The VDisk-to-host mapping was not created because the maximum number of mappings has been reached.

Explanation: The host map was not created because the maximum number of mappings has been reached.

User response: Not applicable.

CMMVC5877E The VDisk-to-host mapping was not created because the maximum number of SCSI LUNs has been allocated.

Explanation: The host map was not created because the maximum number of SCSI LUNs has been allocated.

User response: Not applicable.

CMMVC5878E The VDisk-to-host mapping was not created because this VDisk is already mapped to this host.

Explanation: The host map was not created because this volume is already mapped to this host.

User response: Specify a different volume, and resubmit the command.

CMMVC5879E The VDisk-to-host mapping was not created because a VDisk is already mapped to this host with this SCSI LUN.

Explanation: The host map was not created because this SCSI LUN is already assigned to another mapping.

User response: Specify a different SCSI LUN, and resubmit the command.

CMMVC5880E The virtual disk was not created because a capacity of zero bytes is not allowed for image mode disks.

Explanation: The host map was not created because the volume has a capacity of zero bytes.

User response: Specify a different volume, and resubmit the command.

CMMVC5881E The FlashCopy mapping was not created because an entity that was specified in the command does not exist.

Explanation: The FlashCopy mapping was not created because an entity that was specified in the command does not exist.

User response: Specify a different entity, and resubmit the command.

CMMVC5882E The FlashCopy mapping was not created because a mapping for the source or target virtual disk (VDisk) already exists.

Explanation: The FlashCopy mapping was not created because a mapping for the source or target volume already exists.

User response: Specify a different source or target volume, and resubmit the command.

CMMVC5883E The FlashCopy mapping was not created because the recovery I/O group is associated with the source or target virtual disk (VDisk).

Explanation: The FlashCopy mapping was not created because the recovery I/O group is associated with the source or target volume.

User response: Specify a different recovery I/O group, and resubmit the command.

CMMVC5884E The FlashCopy mapping was not created because the source or target virtual disk (VDisk) cannot be a member of a Remote Copy mapping.

Explanation: The FlashCopy mapping was not created because the source or target volume cannot be a member of a Remote Copy mapping.

User response: Specify a different source or target volume, and resubmit the command.

CMMVC5885E The FlashCopy mapping was not created because this source or target virtual disk (VDisk) cannot be a member of a FlashCopy mapping.

Explanation: The FlashCopy mapping was not created because this source or target volume cannot be a member of a FlashCopy mapping.

User response: Specify a different source or target volume, and resubmit the command.

CMMVC5886E The FlashCopy mapping was not created because the source or target virtual disk (VDisk) is associated with the recovery I/O group.

Explanation: The FlashCopy mapping was not created because the source or target volume is associated with the recovery I/O group.

User response: Specify a different source or target volume, and resubmit the command.

CMMVC5887E The FlashCopy mapping was not created because the source or target virtual disk (VDisk) must not be in router mode.

Explanation: The FlashCopy mapping was not created because the source or target volume must not be in router mode.

User response: Specify a different source or target volume, and resubmit the command.

CMMVC5888E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

User response: Specify the correct entity, and resubmit the command.

CMMVC5889E The FlashCopy mapping was not deleted because an entity that was specified in the command does not exist.

Explanation: The FlashCopy mapping was not deleted because an entity that was specified in the command does not exist.

User response: Specify a different entity, and resubmit the command.

CMMVC5890E The FlashCopy mapping or consistency group was not started because starting consistency group 0 is not a valid operation.

Explanation: The FlashCopy mapping or consistency group was not started because starting consistency group 0 is not a valid operation.

User response: Not applicable.

CMMVC5891E The FlashCopy consistency group was not created because the name is not valid.

Explanation: The FlashCopy consistency group was not created because the name is not valid.

User response: Specify a different name, and resubmit the command.

CMMVC5892E The FlashCopy consistency group was not created because it already exists.

Explanation: The FlashCopy consistency group was not created because it already exists.

User response: Not applicable.

CMMVC5893E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

User response: Specify the correct entity, and resubmit the command.

CMMVC5894E The FlashCopy consistency group was not deleted because you are trying to delete consistency group 0 or the name of the consistency group is not valid.

Explanation: The FlashCopy consistency group was not deleted because the name of the consistency group is not valid or you are trying to delete consistency group 0.

User response: Specify the correct consistency group, and resubmit the command.

CMMVC5895E The FlashCopy consistency group was not deleted because it contains mappings. To delete this consistency group, a forced deletion is required.

Explanation: The FlashCopy consistency group was not deleted because it contains mappings.

User response: Specify that -force option to delete the consistency group.

CMMVC5896E The FlashCopy mapping was not deleted because the mapping or consistency group is in the preparing state. The mapping or consistency group must be stopped first.

Explanation: The FlashCopy mapping was not deleted because the mapping or consistency group is in the preparing state. The mapping or consistency group must be stopped first.

User response: Stop the consistency group, and resubmit the command.

CMMVC5897E The FlashCopy mapping was not deleted because the mapping or consistency group is in the prepared state. The mapping or consistency group must be stopped first.

Explanation: The FlashCopy mapping was not deleted because the mapping or consistency group is in the prepared state. The mapping or consistency group must be stopped first.

User response: Stop the consistency group, and resubmit the command.

CMMVC5898E The FlashCopy mapping was not deleted because the mapping or consistency group is in the copying state. The mapping or consistency group must be stopped first.

Explanation: The FlashCopy mapping was not deleted because the mapping or consistency group is in the copying state. The mapping or consistency group must be stopped first.

User response: Stop the consistency group, and resubmit the command.

CMMVC5899E The FlashCopy mapping was not deleted because the mapping or consistency group is in the stopped state. To delete the mapping, a forced deletion is required.

Explanation: The FlashCopy mapping was not deleted because the mapping or consistency group is in the stopped state.

User response: Specify the -force option to delete the mapping.

CMMVC5900E The FlashCopy mapping was not deleted because the mapping or consistency group is in the suspended state. The mapping or consistency group must be stopped first.

Explanation: The FlashCopy mapping was not deleted because the mapping or consistency group is in the suspended state. The mapping or consistency group must be stopped first.

User response: Stop the consistency group, and resubmit the command.

CMMVC5901E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the preparing state.

Explanation: The FlashCopy mapping was not prepared because the mapping or consistency group is already in the preparing state.

User response: Not applicable.

CMMVC5902E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the prepared state.

Explanation: The FlashCopy mapping was not prepared because the mapping or consistency group is already in the prepared state.

User response: Not applicable.

CMMVC5903E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the copying state.

Explanation: The FlashCopy mapping was not prepared because the mapping or consistency group is already in the copying state.

User response: Not applicable.

CMMVC5904E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the suspended state.

Explanation: The FlashCopy mapping was not prepared because the mapping or consistency group is already in the suspended state.

User response: Not applicable.

CMMVC5905E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the idle state. The mapping or consistency group must be prepared first.

Explanation: The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the idle state.

User response: Prepare the mapping or consistency group, and resubmit the command.

CMMVC5906E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the preparing state.

Explanation: The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the preparing state.

User response: Not applicable.

CMMVC5907E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is already in the copying state.

Explanation: The FlashCopy mapping or consistency group was not started because the mapping or consistency group is already in the copying state.

User response: Not applicable.

CMMVC5908E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopped state. The mapping or consistency group must be prepared first.

Explanation: The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopped state.

User response: Prepare the mapping or consistency group, and resubmit the command.

CMMVC5909E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the suspended state.

Explanation: The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the suspended state.

User response: Not applicable.

CMMVC5910E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the idle state.

Explanation: The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the idle state.

User response: Not applicable.

CMMVC5911E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the preparing state.

Explanation: The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the preparing state.

User response: Not applicable.

CMMVC5912E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopped state.

Explanation: The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopped state.

User response: Not applicable.

CMMVC5913E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the preparing state.

Explanation: The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the preparing state.

User response: Not applicable.

CMMVC5914E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the prepared state.

Explanation: The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the prepared state.

User response: Not applicable.

CMMVC5915E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the copying state.

Explanation: The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the copying state.

User response: Not applicable.

CMMVC5916E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the suspended state.

Explanation: The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the suspended state.

User response: Not applicable.

CMMVC5917E The FlashCopy mapping was not created because there is no memory in which to create the bitmap.

Explanation: The FlashCopy mapping was not created because there is no memory to create the bitmap.

User response: Not applicable.

CMMVC5918E The FlashCopy mapping was not prepared, either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings or unrecovered Global Mirror or Metro Mirror relationships in the I/O group.

Explanation: This error might be caused by a temporary loss of all of the nodes in the I/O group,

which causes all of the FlashCopy mappings and Global and Metro Mirror relationships of the I/O group to be unusable.

User response: Perform the following steps:

1. Ensure that at least one of the nodes in the I/O group of the mapping is online.
2. Fix all of the unfixed events in the event log.
3. Follow the Directed Maintenance Procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5919E The FlashCopy mapping or consistency group was not started, either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings or unrecovered Global Mirror or Metro Mirror relationships in the I/O group.

Explanation: This error might be caused by a temporary loss of all of the nodes in the I/O group, which causes all of the FlashCopy mappings and Global and Metro Mirror relationships of the I/O group to be unusable.

User response: Perform the following steps:

1. Ensure that at least one of the nodes in the I/O group of the mapping is online.
2. Fix all of the unfixed events in the event log.
3. Follow the Directed Maintenance Procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5920E The FlashCopy mapping was not created because the consistency group is not idle.

Explanation: The FlashCopy mapping was not created because the consistency group is not idle.

User response: Not applicable.

CMMVC5921E The properties of the FlashCopy mapping were not modified because the consistency group is not idle.

Explanation: The properties of the FlashCopy mapping were not modified because the consistency group is not idle.

User response: Not applicable.

CMMVC5922E The FlashCopy mapping was not created because the destination virtual disk (VDisk) is too small.

Explanation: The FlashCopy mapping was not created because the destination volume is too small.

User response: Specify a different volume, and resubmit the command.

CMMVC5923E The FlashCopy mapping cannot be created, either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings or unrecovered Global Mirror or Metro Mirror relationships in the I/O group.

Explanation: This error might be caused by a temporary loss of all of the nodes in the I/O group, which causes all of the FlashCopy mappings and Global and Metro Mirror relationships of the I/O group to be unusable.

User response: Perform the following steps:

1. Ensure that at least one of the nodes in the I/O group of the mapping is online.
2. Fix all of the unfixed events in the event log.
3. Follow the Directed Maintenance Procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5924E The FlashCopy mapping was not created because the source and target virtual disks (VDisks) are different sizes.

Explanation: The FlashCopy mapping was not created because the source and target volumes are different sizes.

User response: Specify a different source and target volume that are the same size, and resubmit the command.

CMMVC5925E The remote cluster partnership was not created because it already exists.

Explanation: The remote cluster partnership was not created because it already exists.

User response: Specify a different remote cluster partnership, and resubmit the command.

CMMVC5926E The remote cluster partnership was not created because there are too many partnerships.

Explanation: The remote cluster partnership was not created because there are too many partnerships.

User response: Not applicable.

CMMVC5927E The action failed because the cluster ID is not valid.

Explanation: The action failed because the cluster ID is not valid.

User response: Specify the correct cluster ID, and resubmit the command.

CMMVC5928E The action failed because the cluster name is a duplicate of another cluster.

Explanation: The action failed because the cluster name is a duplicate of another cluster.

User response: Specify a different cluster name, and resubmit the command.

CMMVC5929E The Remote Copy partnership was not deleted because it has already been deleted.

Explanation: The Remote Copy partnership was not deleted because it has already been deleted.

User response: Not applicable.

CMMVC5930E The Remote Copy relationship was not created because an object that was specified in the command does not exist.

Explanation: The Remote Copy relationship was not created because an object that was specified in the command does not exist.

User response: Specify the correct object, and resubmit the command.

CMMVC5931E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is locked.

Explanation: The Remote Copy relationship was not created because the master or auxiliary volume is locked.

User response: Unlock the master or auxiliary volume, and resubmit the command.

CMMVC5932E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a FlashCopy mapping.

Explanation: The Remote Copy relationship was not created because the master or auxiliary volume is a member of a FlashCopy mapping.

User response: Not applicable.

CMMVC5933E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is in the recovery I/O group.

Explanation: The Remote Copy relationship was not created because the master or auxiliary volume is in the recovery I/O group.

User response: Not applicable.

CMMVC5934E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is in the router mode.

Explanation: The Remote Copy relationship was not created because the master or auxiliary volume is in the router mode.

User response: Not applicable.

CMMVC5935E The action failed because an object that was specified in the command does not exist.

Explanation: The action failed because an object that was specified in the command does not exist.

User response: Specify the correct object, and resubmit the command.

CMMVC5936E The action failed because an object that was specified in the command does not exist.

Explanation: The action failed because an object that was specified in the command does not exist.

User response: Specify the correct object, and resubmit the command.

CMMVC5937E The action failed because an object that was specified in the command does not exist.

Explanation: The action failed because an object that was specified in the command does not exist.

User response: Specify the correct object, and resubmit the command.

CMMVC5938E The Remote Copy consistency group was not deleted because the consistency group contains relationships. To delete the consistency group, the force option is required.

Explanation: Remote Copy consistency group was not deleted because the consistency group contains relationships.

User response: Specify the -force option to delete the consistency group.

CMMVC5939E The action failed because the cluster is not in a stable state.

Explanation: The action failed because the cluster is not in a stable state.

User response: Not applicable.

CMMVC5940E The cluster that contains the auxiliary virtual disk (VDisk) is unknown.

Explanation: The cluster that contains the auxiliary volume is unknown.

User response: Not applicable.

CMMVC5941E The cluster that contains the master virtual disk (VDisk) has too many consistency groups.

Explanation: The cluster that contains the master volume has too many consistency groups.

User response: Not applicable.

CMMVC5942E The cluster that contains the auxiliary virtual disk (VDisk) has too many consistency groups.

Explanation: The cluster that contains the auxiliary volume has too many consistency groups.

User response: Not applicable.

CMMVC5943E The specified relationship is not valid.

Explanation: The specified relationship is not valid.

User response: Specify the correct relationship, and resubmit the command.

CMMVC5944E The specified consistency group is not valid.

Explanation: The specified consistency group is not valid.

User response: Specify the correct consistency group, and resubmit the command.

CMMVC5945E The specified master cluster is not valid.

Explanation: The specified master cluster is not valid.

User response: Specify the correct master cluster, and resubmit the command.

CMMVC5946E The specified auxiliary cluster is not valid.

Explanation: The specified auxiliary cluster is not valid.

User response: Specify the correct auxiliary cluster, and resubmit the command.

CMMVC5947E The specified master virtual disk (VDisk) is not valid.

Explanation: The specified master volume is not valid.

User response: Specify the correct master volume, and resubmit the command.

CMMVC5948E The specified auxiliary virtual disk (VDisk) is not valid.

Explanation: The specified auxiliary volume is not valid.

User response: Specify the auxiliary volume, and resubmit the command.

CMMVC5949E The specified relationship is unknown.

Explanation: The specified relationship is unknown.

User response: Specify a different relationship, and resubmit the command.

CMMVC5950E The specified consistency group is unknown.

Explanation: The specified consistency group is unknown.

User response: Specify a different consistency group, and resubmit the command.

CMMVC5951E The operation cannot be performed because the relationship is not a stand-alone relationship.

Explanation: The operation cannot be performed because the relationship is not a stand-alone one.

User response: Not applicable.

CMMVC5952E The relationship and consistency group have different master clusters.

Explanation: The relationship and consistency group have different master clusters.

User response: Not applicable.

CMMVC5953E The relationship and group have different auxiliary clusters.

Explanation: The relationship and group have different auxiliary clusters.

User response: Not applicable.

CMMVC5954E The master and auxiliary virtual disks (VDisks) are different sizes.

Explanation: The master and auxiliary volumes are different sizes.

User response: Not applicable.

CMMVC5955E The maximum number of relationships has been reached.

Explanation: The maximum number of relationships has been reached.

User response: Not applicable.

CMMVC5956E The maximum number of consistency groups has been reached.

Explanation: The maximum number of consistency groups has been reached.

User response: Not applicable.

CMMVC5957E The master virtual disk (VDisk) is already in a relationship.

Explanation: The master volume is already in a relationship.

User response: Specify a different master volume, and resubmit the command.

CMMVC5958E The auxiliary virtual disk (VDisk) is already in a relationship.

Explanation: The auxiliary volume is already in a relationship.

User response: Specify a different auxiliary volume, and resubmit the command.

CMMVC5959E There is a relationship that already has this name on the master cluster.

Explanation: There is a relationship that already has this name on the master cluster.

CMMVC5960E • CMMVC5971E

User response: Specify a different name, and resubmit the command.

CMMVC5960E There is a relationship that already has this name on the auxiliary cluster.

Explanation: There is a relationship that already has this name on the auxiliary cluster.

User response: Specify a different name, and resubmit the command.

CMMVC5961E There is a consistency group that already has this name on the master cluster.

Explanation: There is a consistency group that already has this name on the master cluster.

User response: Specify a different name, and resubmit the command.

CMMVC5962E There is a consistency group that already has this name on the auxiliary cluster.

Explanation: There is a consistency group that already has this name on the auxiliary cluster.

User response: Specify a different name, and resubmit the command.

CMMVC5963E No direction has been defined.

Explanation: No direction has been defined.

User response: Not applicable.

CMMVC5964E The copy priority is not valid.

Explanation: The copy priority is not valid.

User response: Not applicable.

CMMVC5965E The virtual disks (VDisks) are in different I/O groups on the local cluster.

Explanation: The volumes are in different I/O groups on the local cluster.

User response: Not applicable.

CMMVC5966E The master virtual disk (VDisk) is unknown.

Explanation: The master volume is unknown.

User response: Specify a different master volume, and resubmit the command.

CMMVC5967E The auxiliary virtual disk (VDisk) is unknown.

Explanation: The auxiliary volume is unknown.

User response: Specify a different auxiliary volume, and resubmit the command.

CMMVC5968E The relationship cannot be added because the states of the relationship and the consistency group do not match.

Explanation: The relationship cannot be added because the states of the relationship and the consistency group do not match.

User response: Not applicable.

CMMVC5969E The Remote Copy relationship was not created, either because there are no online nodes in the I/O group or because there are unrecovered FlashCopy mappings or unrecovered Global Mirror or Metro Mirror relationships in the I/O group.

Explanation: This error might be caused by a temporary loss of all of the nodes in the I/O group, which causes all of the FlashCopy mappings and Global and Metro Mirror relationships of the I/O group to be unusable.

User response: Perform the following steps:

1. Ensure that at least one of the nodes in the I/O group is online.
2. Fix all of the unfixed events in the event log.
3. Follow the Directed Maintenance Procedures.

You might be required to delete and re-add ALL of the FlashCopy maps and Global and Metro Mirror relationships in the I/O group.

Resubmit the command.

CMMVC5970E The Remote Copy relationship was not created because there is not enough memory.

Explanation: The Remote Copy relationship was not created because there is not enough memory.

User response: Not applicable.

CMMVC5971E The operation was not performed because the consistency group contains no relationships.

Explanation: The operation was not performed because the consistency group contains no relationships.

User response: Not applicable.

CMMVC5972E The operation was not performed because the consistency group contains relationships.

Explanation: The operation was not performed because the consistency group contains relationships.

User response: Not applicable.

CMMVC5973E The operation was not performed because the consistency group is not synchronized.

Explanation: The operation was not performed because the consistency group is not synchronized.

User response: Specify the Force option when starting the consistency group.

CMMVC5974E The operation was not performed because the consistency group is offline.

Explanation: The operation was not performed because the consistency group is offline.

User response: Not applicable.

CMMVC5975E The operation was not performed because the cluster partnership is not connected.

Explanation: The operation was not performed because the cluster partnership is not connected.

User response: Not applicable.

CMMVC5976E The operation was not performed because the consistency group is in the freezing state.

Explanation: The operation was not performed because the consistency group is in the freezing state.

User response: Not applicable.

CMMVC5977E The operation was not performed because it is not valid given the current consistency group state.

Explanation: The operation was not performed because it is not valid given the current consistency group state.

User response: Not applicable.

CMMVC5978E The operation was not performed because the relationship is consistent but is not synchronized. Restarting the relationship by using the -force parameter will make the relationship inconsistent until the background copy has completed.

Explanation: Input transactions have occurred on either the primary or secondary volumes since the ConsistentStopped or Idling state has occurred. Because the relationship is no longer synchronized, the state of the relationship is now Stopped.

The -force parameter of the **svctask startrelationship** command is required when the relationship is not synchronized because consistency would be lost by starting the copy operation. Submitting the **svctask startrelationship** command on an unsynchronized relationship without using the -force parameter is not supported.

If a relationship is in the InconsistentStopped, InconsistentCopying or ConsistentSynchronized state, the -force parameter is not required, but is supported.

User response: Consider using the -force parameter of the **svctask startrelationship** command, if appropriate.

CMMVC5980E The operation was not performed because the master and auxiliary clusters are not connected.

Explanation: The operation was not performed because the master and auxiliary clusters are not connected.

User response: Not applicable.

CMMVC5981E The operation was not performed because the relationship is in the freezing state.

Explanation: The operation was not performed because the relationship is in the freezing state.

User response: Not applicable.

CMMVC5982E The operation was not performed because it is not valid given the current relationship state.

Explanation: The operation was not performed because it is not valid given the current relationship state.

User response: Not applicable.

CMMVC5983E dump file was not created. This may be due to the file system being full.

Explanation: dump file was not created. This may be due to the file system being full.

User response: Not applicable.

CMMVC5984E The dump file was not written to disk. This may be due to the file system being full.

Explanation: The dump file was not written to disk. This may be due to the file system being full.

User response: Not applicable.

CMMVC5985E The action failed because the specified directory is not permitted for this command.

Explanation: You have attempted to copy, delete, or list dumps from a directory that is not valid. A list of valid directories for these commands is provided in the documentation.

User response: Ensure that the directory you specify is valid, and resubmit the command.

CMMVC5986E The tracing of I/O operations was not started because the virtual disk (VDisk) or managed disk (MDisk) failed to return any statistics.

Explanation: The tracing of I/O operations was not started because the volume or managed disk (MDisk) failed to return statistics.

User response: Not applicable.

CMMVC5987E *VALUE* is not a valid command line option.

Explanation: The specified string that you have entered is not a supported command line option.

User response: Specify a supported option, and resubmit the command.

CMMVC5988E command should not be run by the root userid. Use the admin userid.

Explanation: This command should not be issued if you are logged in with a root user ID. Use the admin userid.

User response: Log off of the root user ID and log in as admin.

CMMVC5989E The operation was not performed because the relationship is offline.

Explanation: The operation was not performed because the relationship is offline.

User response: Not applicable.

CMMVC5990E The FlashCopy consistency group was not stopped as there are no FlashCopy mappings within the group.

Explanation: The FlashCopy consistency group was not stopped as there are no FlashCopy mappings within the group.

User response: Not applicable.

CMMVC5991E The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.

Explanation: The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.

User response: Not applicable.

CMMVC5992E The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.

Explanation: The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.

User response: Not applicable.

CMMVC5993E The specified upgrade package does not exist.

Explanation: The specified upgrade package does not exist.

User response: Not applicable.

CMMVC5994E Error in verifying the signature of the upgrade package.

Explanation: The system could not verify the signature of the upgrade package due to the following reasons:

- There is not enough space on the system to copy the file.
- The package is incomplete or contains errors.

User response: If the copy failed with an error indicating that there was insufficient space on the system, free up additional space on your system. Otherwise, ensure that the cluster time and date stamp on the signature is correct. (For example, the time and date cannot be in the future.)

CMMVC5995E Error in unpacking the upgrade package.

Explanation: The most likely cause of this error is lack of system space.

User response: Reboot the node and unpack the upgrade package again.

CMMVC5996E The specific upgrade package cannot be installed over the current version.

Explanation: The upgrade package is not compatible with the current version or the system.

User response: Check the available upgrade packages and find the correct upgrade package for your current version and for your system. If the upgrade package is correct for your system, check the version requirements for the package. You might have to upgrade the current version to an intermediate version before you upgrade to the latest version. (For example, if your current version is 1 and you are trying to upgrade to version 3, you might need to upgrade to version 2 before applying the version 3 upgrade.)

CMMVC5999W Featurization for this facility has not been enabled.

Explanation: Featurization for this facility has not been enabled.

User response: Not applicable.

CMMVC6000W Featurization for this facility has not been enabled.

Explanation: Featurization for this facility has not been enabled.

User response: Not applicable.

CMMVC6001E The FlashCopy consistency group was not started as there are no FlashCopy mappings within the group.

Explanation: The FlashCopy consistency group was not started as there are no FlashCopy mappings within the group.

User response: Create a FlashCopy within the appropriate group.

CMMVC6002E This command can only be run on a node that is in the service state.

Explanation: This command can only be run on a node that is in the service state.

User response: Not applicable.

CMMVC6003E This command can not be run on a node that is in the service state.

Explanation: This command can not be run on a node that is in the service state.

User response: Not applicable.

CMMVC6004E The delimiter value *VALUE* is invalid.

Explanation: The specified value is not a valid delimiter value.

User response: Specify a different delimiter.

CMMVC6005E The view request failed as the specified object is not a member of an appropriate group.

Explanation: A view was request on an object that has been incorrectly initialized.

User response: Ensure that the object is correctly initialized before resubmitting the view request.

CMMVC6006E The managed disk (MDisk) was not deleted because the resource was busy.

Explanation: An attempt was made to delete an MDisk from a storage pool that is being used as a source and destination for migration operations.

User response: Ensure that the storage pool is not being used for migration operations before reissuing the command.

CMMVC6007E The two passwords that were entered do not match.

Explanation: The two passwords entered for verification of your password change were not the same.

User response: Re-enter the passwords.

CMMVC6008E The key already exists.

Explanation: An attempt was made to load a duplicate SSH key.

User response: Not applicable.

CMMVC6009E Unable to malloc a block of memory in which to copy the returned data.

Explanation: The command line was unable to allocate a block of memory in which to copy the results of the query.

User response: Resubmit the command. If the problem persists, contact IBM technical support for assistance.

CMMVC6010E Unable to complete the command as there are insufficient free extents, or the command requested an expansion of 0 size.

Explanation: There are not enough free extents to meet the request.

User response: Not applicable.

CMMVC6011E This cluster is part of a remote cluster partnership. Because this upgrade package will make changes to the cluster state, it cannot be applied to the current code level until all remote cluster partnerships are deleted.

Explanation: You have attempted to apply software when a Remote Copy relationship to a remote cluster exists.

User response: Delete the Remote Copy relationship to the remote clusters, and resubmit the command.

CMMVC6012W The virtualized storage capacity is approaching the amount that you are licensed to use.

Explanation: The requested action has completed. However, the limits permitted by the license you purchased are approaching.

User response: Subsequent actions might require that you increase your licensed limits.

CMMVC6013E The command failed because there is a consistency group mismatch on the aux cluster.

Explanation: The action has failed as there was a difference in attributes between the Metro Mirror consistency groups involved.

User response: Ensure that the attributes of the two Metro Mirror consistency groups match before resubmitting the command.

CMMVC6014E The command failed because the requested object is either unavailable or does not exist.

Explanation: The command failed because the requested object is either unavailable or does not exist.

User response: Ensure that all parameters have been correctly entered. If this is the case the determine why the object is unavailable, then resubmit the command.

CMMVC6015E A delete request is already in progress for this object.

Explanation: A delete request is already in progress for this object.

User response: Not applicable.

CMMVC6016E The action failed as there would be, or are, no more disks in the MDisk group.

Explanation: The action failed as there would be, or are, no more disks in the I/O group.

User response: Ensure that all parameters have been correctly entered.

CMMVC6017E A parameter or argument contains invalid characters. Ensure that all characters are ASCII.

Explanation: The command-line interface (CLI) will only accept ASCII input.

User response: Ensure that all input to the CLI is ASCII, then resubmit the command.

CMMVC6018E The software upgrade pre-install process failed.

Explanation: The software upgrade failed as there was an error during the preprocessing. The package is either not valid or corrupted.

User response: Ensure the package is a valid upgrade package. Download the package from the source location again as it might have been corrupted during a network transfer.

CMMVC6019E The software upgrade failed as a node pended as the upgrade was in progress.

Explanation: The software upgrade failed as a node pended as the upgrade was in progress.

User response: Ensure that all nodes are online and available before restarting the upgrade process.

CMMVC6020E The software upgrade failed as the system was unable to distribute the software package to all nodes.

Explanation: The software upgrade failed as the system was unable to distribute the software package to all nodes.

User response: Ensure that all nodes are correctly zoned and that all nodes are online and can see the other nodes in the cluster. You might also want to check the event log.

CMMVC6021E The system is currently busy performing another request. Try again later.

Explanation: The requested action failed as the system is processing another request.

User response: Wait before resubmitting the request.

CMMVC6022E The system is currently busy performing another request. Try again later.

Explanation: The requested action failed as the system is processing another request.

User response: Wait before resubmitting the request.

CMMVC6023E The system is currently busy performing another request. Try again later.

Explanation: The requested action failed as the system is processing another request.

User response: Wait before resubmitting the request.

CMMVC6024E The auxiliary VDisk entered is invalid.

Explanation: The auxiliary volume is entered as a parameter in the command-line interface is not a valid auxiliary volume.

User response: Select a valid auxiliary volume, and resubmit the command.

CMMVC6025E The RC consistency group Master cluster is not the local cluster.

Explanation: The auxiliary volume is entered as a parameter in the command-line interface is not a valid auxiliary volume.

User response: Resubmit the command with a consistency group that belongs to the local cluster.

CMMVC6026E The RC consistency group is not in the stopped state.

Explanation: The action failed as the Metro Mirror consistency group is not in the stopped state.

User response: Ensure that the Metro Mirror consistency group is in the stopped state before resubmitting the command.

CMMVC6027E The RC consistency group is not the primary master.

Explanation: The RC consistency group requested in the command is not the Metro Mirror primary master.

User response: Ensure that the parameters have been entered correctly on the command line.

CMMVC6028E This upgrade package cannot be applied to the current software level because it contains changes to the cluster state and there are remote cluster partnership defined.

Explanation: The action failed because there is a connected remote cluster. The upgrade cannot be applied because it would render the remote cluster at a different code level to the remote cluster.

User response: Ensure that the cluster partnership is unconfigured before resubmitting the command. Ensure

that you unconfigure the remote cluster and upgrade the code on it before reconfiguring the cluster partnership.

CMMVC6029E All nodes must have identical code level before a concurrent code upgrade can be performed.

Explanation: The concurrent upgrade failed as two or more nodes were at differing code levels. All nodes must be at the same code level before a software upgrade can be performed.

User response: Use the service assistant to bring all nodes to the same level before resubmitting the concurrent upgrade.

CMMVC6030E The operation was not performed because the FlashCopy mapping is part of a consistency group. The action must be performed at the consistency group level.

Explanation: An attempt was made to stop a FlashCopy mapping. This failed as the FlashCopy mapping is part of a consistency group.

User response: Issue the stop command to the FlashCopy consistency group. This will stop all FlashCopies within that group that are in progress.

CMMVC6031E The operation was not performed because the FlashCopy consistency group is empty.

Explanation: An attempt was made to prestart an empty FlashCopy consistency group.

User response: Not applicable.

CMMVC6032E The operation was not performed because one or more of the entered parameters is invalid for this operation.

Explanation: An parameter that is not valid was entered for the command.

User response: If attempting to change the I/O group to which the volume belongs, ensure that the volume is not already a part of the group.

CMMVC6033E The action failed due to an internal error.

Explanation: An internal error caused the action to fail.

User response: Not applicable.

CMMVC6034E The action failed because the maximum number of objects has been reached.

Explanation: The action failed because the maximum number of objects has been reached.

User response: Not applicable.

CMMVC6035E The action failed as the object already exists.

Explanation: An operation was requested to create an object that already exists.

User response: Ensure that the name you are attempting to apply to a new object does not exist, or change the name before re-issuing the command.

CMMVC6036E An invalid action was requested.

Explanation: The action failed because it is not a valid action with the command that was issued.

User response: Issue an action that is valid with the command.

CMMVC6037E The action failed as the object is not empty.

Explanation: The action failed because an object was specified.

User response: Resubmit the command without specifying an object.

CMMVC6038E The action failed as the object is empty.

Explanation: The action failed because an object was not specified.

User response: Specify an object, and resubmit the command.

CMMVC6039E The action failed as the object is not a member of a group.

Explanation: The action failed because the object is not a member of a group.

User response: Specify an object that is part of a group, and resubmit the command.

CMMVC6040E The action failed as the object is not a parent.

Explanation: The action failed because the object is not a parent object.

User response: Specify an object that is a parent, and resubmit the command.

CMMVC6041E The action failed as the cluster is full.

Explanation: The action failed because the cluster is full.

User response: Remove data from the cluster, and resubmit the command.

CMMVC6042E The action failed as the object is not a cluster member.

Explanation: The action failed because the object is not a member of the cluster.

User response: Specify an object that is a member of the cluster, and resubmit the command.

CMMVC6043E The action failed as the object is a member of a group.

Explanation: The action failed because the object is a member of a group.

User response: Specify an object that is not a member of a group, and resubmit the command.

CMMVC6044E The action failed as the object is a parent.

Explanation: The action failed because the object is a parent object.

User response: Specify an object that is not a parent object, and resubmit the command.

CMMVC6045E The action failed, as the -force flag was not entered.

Explanation: The action failed because the -force option was not entered.

User response: Specify the -force option in the command.

CMMVC6046E The action failed as too many candidates were selected.

Explanation: The action failed because too many candidates were specified.

User response: Specify fewer candidates in the command.

CMMVC6047E The action failed as too few candidates were selected.

Explanation: An action was requested with too few candidate objects.

User response: Determine the correct number of candidates required for the specific command and reissue the command.

CMMVC6048E The action failed as the object is busy.

Explanation: The action failed because the object is busy.

User response: Not applicable.

CMMVC6049E The action failed as the object is not ready.

Explanation: The action failed because the object is not ready.

User response: Not applicable.

CMMVC6050E The action failed as the command was busy.

Explanation: The action failed because the command is busy.

User response: Not applicable.

CMMVC6051E An unsupported action was selected.

Explanation: The action failed because it is not valid with the command.

User response: Specify an action that is valid with the command.

CMMVC6052E The action failed as the object is a member of a FlashCopy mapping.

Explanation: The object is a member of a FlashCopy mapping, thus it cannot be deleted.

User response: Specify an object that is not a member of a FlashCopy mapping, or remove the object from the FlashCopy mapping.

CMMVC6053E An invalid WWPN was entered.

Explanation: A worldwide port name (WWPN) that is not valid was specified.

User response: Specify a valid WWPN.

CMMVC6054E The action failed as not all nodes are online.

Explanation: One or more nodes are not online.

User response: Check that each node is online, and resubmit the command.

CMMVC6055E The action failed as an upgrade is in progress.

Explanation: The action failed because a software upgrade is in progress.

User response: Wait for the software upgrade to

complete, and resubmit the command.

CMMVC6056E The action failed as the object is too small.

Explanation: The action failed because the object is too small.

User response: Specify a different object, and resubmit the command.

CMMVC6057E The action failed as the object is the target of a FlashCopy mapping.

Explanation: The object is the target of a FlashCopy mapping, thus it cannot be deleted.

User response: Specify an object that is not the target of a FlashCopy mapping, or remove the object from the FlashCopy mapping.

CMMVC6058E The action failed as the object is in the recovery HWS.

Explanation: An attempt was made to perform an operation on a node that is in the recovery I/O group.

User response: Get the node into one of the other I/O groups and reissue the command.

CMMVC6059E The action failed as the object is in an invalid mode.

Explanation: The action failed because the object is in the wrong mode.

User response: Check that the object is in the correct mode, and resubmit the command.

CMMVC6060E The action failed as the object is being deleted.

Explanation: The action failed because the object is being deleted.

User response: Not applicable.

CMMVC6061E The action failed as the object is being resized.

Explanation: The action failed because the object is being resized.

User response: Check that the object is in the correct mode, and resubmit the command.

CMMVC6062E The action failed as the object is being moved between HWS.

Explanation: An attempt was made to perform an action against an object that is currently being moved between I/O groups.

CMMVC6063E • CMMVC6077E

User response: Re-issue the command when the move operation has completed.

CMMVC6063E The action failed as there are no more disks in the group.

Explanation: An attempt was made to perform an action against a group that contained no disks.

User response: Either add disks to the group and reissue the command, or select another group against which to execute the action.

CMMVC6064E The action failed as the object has an invalid name.

Explanation: An attempt was made to create or rename an object using a name that is not valid.

User response: Use a name that meets the naming standards and reissue the command.

CMMVC6065E The action failed as the object is not in a group.

Explanation: An attempt was made to perform an action on an object that was not in an appropriate group.

User response: Ensure that the object is a member of an appropriate group and reissue the command.

CMMVC6066E The action failed as the system is running low on memory.

Explanation: The system is running low on memory.

User response: Not applicable.

CMMVC6067E The action failed as the SSH key was not found.

Explanation: An attempt was made to perform an action using an SSH key that does not exist.

User response: Re-issue the command using a key that does exist.

CMMVC6068E The action failed as there are no free SSH keys.

Explanation: An attempt was made to use an SSH key when there are no free SSH keys.

User response: Upload additional keys and reissue the command.

CMMVC6069E The action failed as the SSH key is already registered.

Explanation: An attempt was made to register an SSH key that was already registered.

User response: Not applicable.

CMMVC6070E An invalid or duplicated parameter, unaccompanied argument, or incorrect argument sequence has been detected. Ensure that the input is as per the help.

Explanation: The parameters entered for a command were not valid.

User response: Correct the parameters and reissue the command.

CMMVC6071E The VDisk-to-host mapping was not created because the VDisk is already mapped to a host.

Explanation: The volume is already mapped to a host.

User response: Not applicable.

CMMVC6073E The maximum number of files has been exceeded.

Explanation: The maximum number of files has been exceeded.

User response: Not applicable.

CMMVC6074E The command failed as the extent has already been assigned.

Explanation: The command failed as the extent has already been assigned.

User response: Assign a different extent, and resubmit the command.

CMMVC6075E The expand failed as the last extent is not a complete extent.

Explanation: The expand failed as the last extent is not a complete extent.

User response: Assign a different extent, and resubmit the command.

CMMVC6076E The command failed because the virtual disk cache is not empty. Either wait for the cache to flush or use the force flag to discard the contents of the cache.

Explanation: The command failed due to an error during the flushing of the volume.

User response: Not applicable.

CMMVC6077E WARNING - Unfixed errors should be fixed before applying software upgrade. Depending on the nature of the errors, they might cause the upgrade process to fail. It is highly recommended to fix these errors before proceeding. If a particular error cannot

be fixed, contact the support center.

Explanation: Unfixed errors should be fixed before applying software upgrade. Depending on the nature of the errors, they might cause the upgrade process to fail. It is highly recommended to fix these errors before proceeding.

User response: If the error cannot be fixed, contact the support center.

CMMVC6078E The action failed because the object is in an invalid mode.

Explanation: An attempt was made to perform an action against an object in a mode that did not allow for that action to be performed.

User response: Get the object into a suitable mode and reissue the command.

CMMVC6079E Metadata recovery could not complete the operation because a parameter is invalid.

Explanation: Metadata recovery could not complete the operation because a parameter is not valid.

User response:

CMMVC6081E Metadata Recovery is busy processing the previous operation.

Explanation: Metadata Recovery is busy processing the previous operation.

User response:

CMMVC6082E The attempt to abort metadata recovery failed because the previous operation has completed.

Explanation: The attempt to cancel metadata recovery failed because the previous operation has completed.

User response: None.

CMMVC6083E Metadata recovery could not find a valid dumpfile required for the rebuild operation.

Explanation: Metadata recovery could not find a valid dumpfile required for the rebuild operation.

User response:

CMMVC6084E Metadata recovery could not create/open/write the scan file, the disk might be full.

Explanation: Metadata recovery could not create/open/write the scan file, the disk might be full.

User response:

CMMVC6085E Metadata recovery could not create/open/write the dump file, the disk might be full.

Explanation: Metadata recovery could not create/open/write the dump file, the disk might be full.

User response:

CMMVC6086E Metadata recovery could not create/open/write the progress file, the disk might be full.

Explanation: Metadata recovery could not create/open/write the progress file, the disk might be full.

User response:

CMMVC6087E Metadata recovery could not map the buffers necessary to complete the operation.

Explanation: Metadata recovery could not map the buffers necessary to complete the operation.

User response:

CMMVC6088E The lba at which metadata recovery was requested does not contain metadata.

Explanation: The lba at which metadata recovery was requested does not contain metadata.

User response:

CMMVC6089E The metadata at the requested lba is flagged as invalid.

Explanation: The metadata at the requested lba is flagged as not valid.

User response:

CMMVC6090E The metadata header checksum verification failed.

Explanation: The metadata header checksum verification failed.

User response:

CMMVC6091E The metadata region checksum verification failed.

Explanation: The metadata region checksum verification failed.

User response:

CMMVC6092E The metadata recovery operation was aborted.

Explanation: The metadata recovery operation was cancelled.

User response:

CMMVC6093E Metadata recovery internal error - (read only)

Explanation: Metadata recovery internal error - (read only)

User response:

CMMVC6095E Metadata recovery encountered the end of the disk.

Explanation: Metadata recovery encountered the end of the disk.

User response:

CMMVC6096E The metadata recovery task could not be initiated because the required back-end resource could not be found.

Explanation: The back-end resource that is required for the task is unavailable.

User response: Ensure that the required back-end resource is available, and reinitiate the task.

CMMVC6097E The metadata recovery task could not be initiated because the system was unable to send the required I/O to the back-end resource.

Explanation: The back-end resource is possibly not configured properly.

User response: Ensure that the required back-end resource is accessible, and reinitiate the task.

CMMVC6098E The copy failed as the specified node is the configuration node.

Explanation: The copy failed because the specified node is the configuration node.

User response: Not applicable.

CMMVC6100E *OPTION* not consistent with *ACTION*

Explanation: The specified option is not supported for the specified action.

User response: Remove the option, and resubmit the command.

CMMVC6101E *OPTION* not consistent with *OPTION*

Explanation: The two specified options cannot be used together.

User response: Remove one of the options, and resubmit the command.

CMMVC6102E *OPTION* and *OPTION* are alternatives

Explanation: The two specified options are alternatives, and cannot be used together.

User response: Remove one of the options, and resubmit the command.

CMMVC6103E Problem with *FILENAME* : *DETAILS*

Explanation: A problem occurred when opening the specified file. Determine the cause of the problem and correct it before trying again.

User response: Correct the problem, and resubmit the command.

CMMVC6104E Action *ACTION* not run

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6105E Different names for source *SOURCE_CLUSTER_NAME* and target *TARGET_CLUSTER_NAME* clusters

Explanation: The backup configuration cannot be restored to the target cluster because the source and target cluster have different names.

User response: Perform one of the following actions: (1) Use a different backup configuration. (2) Delete the cluster and recreate it with the same name as that stored in the backup configuration file.

CMMVC6106W Target cluster has non-default *id_alias* *ALIAS* .

Explanation: The specified *id_alias* of the target cluster is a non-default value. Clusters should have the default value. The non-default value suggests that the cluster is customized and is not suitable for restoration. Restoration changes the *id_alias*.

User response: Change the *id_alias* to a default value, and resubmit the command.

CMMVC6107E *NUMBER_OF_OBJECTS* *io_grp* objects in target cluster; *NUMBER_OF_REQUIRED_OBJECTS* are required

Explanation: The number of I/O groups in the target cluster is not sufficient to accommodate the I/O groups defined in the backup configuration file. Determine why there are not enough I/O groups.

User response: Correct the problem, and resubmit the command.

CMMVC6108I Disk controller system with a WWNN of *WWNN_VALUE* found.

Explanation: A disk controller system with the required WWNN has been found.

User response: Not applicable.

CMMVC6109E Disk controller system with a WWNN of *WWNN_VALUE* not available.

Explanation: A disk controller system with the specified WWNN has been found. Ensure that the specified disk controller system is available to the cluster.

User response: Ensure that the required disk controller system is available to the cluster, and resubmit the command.

CMMVC6110E Bad code level: *VALUE* .

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6111E The cluster code level could not be determined from *VALUE* .

Explanation: The code level of the cluster could not be determined. The code level should be of the format *x.y.z*, where *x*, *y*, and *z* are integers.

User response: If the cause of the problem cannot be determined, contact IBM technical support for assistance.

CMMVC6112W *OBJECT_TYPE* *OBJECT_NAME* has a default name.

Explanation: An object in the cluster has a default name. This can cause problems when restoring a cluster because default names are changed during restoration. Object IDs are also changed during restoration.

User response: Choose an appropriate name for each object in the cluster, and resubmit the command.

CMMVC6113E The command *COMMAND* has failed with return code *RETURN_CODE* .

Explanation: An attempt to run a command remotely failed using secure communications.

User response: Determine the cause of the problem, and resubmit the command.

CMMVC6114E No help for action *ACTION* .

Explanation: There is no help for the specified action topic.

User response: Not applicable.

CMMVC6115W Feature *FEATURE_PROPERTY* mismatch: *VALUE* expected; *VALUE* found.

Explanation: The features in the backup configuration file and the target cluster do not match. There should be an exact match between the two. Nevertheless, the restore of the configuration can continue.

User response: Not applicable.

CMMVC6116I Feature match for *FEATURE* .

Explanation: The features in the backup configuration file and the target cluster are an exact match.

User response: Not applicable.

CMMVC6117E *FIX_OR_FEATURE* is not available.

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6118I *TYPE* with *PROPERTY* *PROPERTY_VALUE* and *PROPERTY* *PROPERTY_VALUE* found.

Explanation: An object in the cluster has been found with the correct properties.

User response: Not applicable.

CMMVC6119E *TYPE* with *PROPERTY* *PROPERTY_VALUE* not found.

Explanation: An object in the cluster with the correct properties has not been found. Restoration cannot proceed without the object.

User response: Determine why the object cannot be found. Ensure that the object is available, and resubmit the command.

CMMVC6120E Target is not the configuration node

Explanation: The target is not the configuration node.

User response: Redirect the action against the configuration node, and resubmit the command.

CMMVC6121E No cluster id or id_alias in backup configuration.

Explanation: Neither the cluster id_alias nor the ID can be extracted from the backup configuration file.

User response: If the cause of the problem cannot be determined, contact IBM technical support for assistance.

CMMVC6122E No TYPE with PROPERTY VALUE is present in the table.

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6123E No PROPERTY for TYPE NAME .

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6124E No TYPE with PROPERTY VALUE

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6125E No unique ID for TYPE NAME

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6126E No TYPE with unique ID VALUE

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6127I The SSH key IDENTIFIER for USER is already defined; the SSH key will not be restored

Explanation: An identical SSH key for this user is already defined on the cluster. Therefore, the key in the backup file will not be restored.

User response: Specify a different SSH key, and resubmit the command.

CMMVC6128W DIRECTORY

Explanation: The files in the specified directory cannot be listed.

User response: Determine why the files cannot be listed, correct the problem, and resubmit the command.

CMMVC6129E VDisk-to-host mapping objects have VDisk_UID values that are not consistent.

Explanation: All of the host mapping objects do not have the same number for the volume LUN instance. Therefore, there is a possibility the backup configuration file is corrupt. The LUN instance number should be the same for all host mapping objects that are associated with a specific volume. The LUN instance number is incorporated into the volume ID property.

User response: Determine why the LUN instance number is not the same, correct the problem, and resubmit the command.

CMMVC6130W Inter-cluster PROPERTY VALUE will not be restored.

Explanation: The restoration of inter-cluster objects is not supported.

User response: Not applicable.

CMMVC6131E No location cluster information

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6132E The object OBJECT of type TYPE has a property PROPERTY with an incorrect value INCORRECT_VALUE . The operation cannot proceed until the property has the correct value CORRECT_VALUE . Take administrative action to change the value and try again.

Explanation: The specified object has the specified property of the specified type with the specified incorrect value. The property most likely reflects the state of the object.

User response: Change the state to the required value, and resubmit the command.

CMMVC6133E Required TYPE property PROPERTY not found

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6134E No argument for *OPTION*

Explanation: No argument has been supplied for the specified option, which requires an argument.

User response: Supply an argument, and resubmit the command.

CMMVC6135E Argument *VALUE* for *OPTION* is not valid.

Explanation: The specified argument that you have supplied is not valid for the specified option.

User response: Supply an valid argument, and resubmit the command.

CMMVC6136W No SSH key file *FILENAME*

Explanation: The specified file, which should contain the SSH key, is not present and will not be restored. The backup operation will continue.

User response: No action is required. You might have to manually restore the key.

CMMVC6137W No SSH key file *FILENAME*; key not restored

Explanation: An SSH key cannot be restored because the specified file, which is expected to contain the SSH key, is not present. The restore operation will continue.

User response: After the restore is complete, locate the file containing the key, and perform one of the following actions: (1) Rename the file so that it has the correct name, and resubmit the command. (2) Restore the key manually using the `svctask addsshkey` command.

CMMVC6138E *OPTION* is required

Explanation: An option is missing. The option might be listed as optional, but circumstances make the option mandatory.

User response: Supply the option, and resubmit the command.

CMMVC6139E Incorrect XML tag nesting in *FILENAME*

Explanation: There is a problem with the content of a configuration file. There is a problem parsing the XML in the file, because the XML records are not consistent. The file might be corrupt, or the file has been truncated.

User response: Replace this copy with a good copy, and resubmit the command. If the problem persists, contact IBM technical support for assistance.

CMMVC6140E No default name for type *TYPE*

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6141E The option *OPTION* does not support an argument.

Explanation: An argument has been supplied for an option that does not support one.

User response: Remove the argument, and resubmit the command.

CMMVC6142E Existing *OBJECT_TYPE* *OBJECT_NAME* has a non-default name.

Explanation: The specified object in the target default cluster has a non-default name. This suggests that the cluster was customized. The cluster is therefore not suitable for restoration.

User response: Reset the cluster as per the instructions for restoring the cluster configuration, and resubmit the command.

CMMVC6143E The required configuration file *FILENAME* does not exist.

Explanation: A file that is critical for successful operation is missing.

User response: Not applicable.

CMMVC6144W The object with default name *NAME* has been restored as *SUBSTITUTE_NAME*.

Explanation: An object with a default name has been restored with a different name. Ensure that you account for this name change when using the restored cluster in the future. To avoid this problem in the future, choose an appropriate name for each object in the cluster.

User response: Choose an appropriate name for each object in the cluster.

CMMVC6145I First use the *COMMAND* -prepare command.

Explanation: This advisory is given prior to CMMVC6103E when an intermediate file is missing.

User response: Not applicable.

CMMVC6146E Problem parsing *OBJECT_TYPE* data: *LINE*

Explanation: An unexpected error has occurred.

User response: Contact the support center.

CMMVC6147E *TYPE NAME* has a name beginning with *PREFIX* .

Explanation: An object has been encountered that has a name beginning with the specified reserved prefix. The only valid reason for an object with this kind of name is that a restoration command did not complete successfully.

User response: Ensure that no object uses the reserved prefix in its name, and resubmit the command.

CMMVC6148E Target cluster has *NUMBER_OF_EXISTING_OBJECTS* objects of type *TYPE* instead of *NUMBER_OF_REQUIRED_OBJECTS* .

Explanation: The target cluster does not have the specified required number of objects of the specified type.

User response: Correct the problem, and resubmit the command.

CMMVC6149E An action is required.

Explanation: An action is required to run the command.

User response: Supply an action, and resubmit the command.

CMMVC6150E The action *ACTION* is not valid.

Explanation: The specified action that you have entered is not valid.

User response: Specify a valid action, and resubmit the command.

CMMVC6151E The option *OPTION* is not valid.

Explanation: The specified option that you have entered is not valid.

User response: Specify a valid option, and resubmit the command.

CMMVC6152E VDisk *VDISK_NAME* instance number *INSTANCE_NUMBER* is not valid.

Explanation: The volume cannot be restored because the instance number, which must be a hexadecimal number, is not valid.

User response: Contact IBM technical support for assistance.

CMMVC6153E *OBJECT* not consistent with *ACTION*

Explanation: The specified object is not supported for the specified action.

User response: Remove the object, and resubmit the command.

CMMVC6154E Required *OBJECT_TYPE* property *PROPERTY_NAME* has a null value.

Explanation: An unexpected error has occurred.

User response: Contact IBM technical support for assistance.

CMMVC6155I The command *COMMAND* processing has completed successfully.

Explanation: Only information and warning messages are issued.

User response: Not applicable.

CMMVC6156W *COMMAND* processing completed with errors.

Explanation: Processing was not successful.

User response: Not applicable.

CMMVC6164E The SVCCONFIG CRON job, which runs overnight on a daily overnight, has failed.

Explanation: The SVCCONFIG CRON job, which runs overnight on a daily overnight, has failed.

User response: Resolve any hardware and configuration problems that you are experiencing on the 2145 cluster. If the problem persists, contact IBM technical support for assistance.

CMMVC6165E The target is not the original configuration node with a WWNN of *WWNN_VALUE* .

Explanation: A backup configuration can only be restored to the original configuration node.

User response: Recreate the default cluster with the correct configuration node, and resubmit the command.

CMMVC6166E The property *PROPERTY* of the object *OBJECT* has changed during *svccfg restore -execute*.

Explanation: The integrity of the restoration cannot be guaranteed.

User response: Resubmit the command from *svccfg restore -prepare*.

CMMVC6181E The target cluster contains an object that has a counterpart in the configuration to be restored, and has the correct ID.

Explanation: The indicated property has an unexpected value.

User response: Check that the correct (matching) backup configuration file (svc.config.backup.xml) is being provided and if it is, use the -force option to ignore the discrepancy. Otherwise, provide the correct file and try again.

CMMVC6182W An object that does not contribute to the fabric of the configuration cannot be restored because its configuration does not permit it to be created.

Explanation: An object that does not contribute to the fabric of the configuration cannot be restored because its configuration does not permit it to be created. For example, a host can only be created if it has at least one port.

User response: Not applicable.

CMMVC6186W The IO group *IO_GROUP_NAME* has been restored with ID *ID_VALUE* instead of *ID_VALUE* .

Explanation: This can occur when the configuration node is different from the node that was used to create the original cluster. This affects the SCSI Inquiry value for the I/O group.

User response: Not applicable.

CMMVC6200E The action failed because of incompatible software.

Explanation: The software version on one or more nodes is incompatible with the new version.

User response: Refer to the compatibility requirements for the software version you are adding. Update the cluster to meet the compatibility requirements, and then perform the upgrade.

CMMVC6201E The node could not be added because of incompatible software. The status code is *STATUS_CODE* .

Explanation: The node could not be added because of incompatible software.

User response: Upgrade the software on the node that has been rejected to the same level of software as the cluster to which it will be added, and resubmit the command.

CMMVC6202E The cluster was not modified because the IP address is not valid.

Explanation: An attempt was made to change the IP address of a cluster to an address that is not valid.

User response: Correct the address and reissue the command.

CMMVC6203E The action failed because the directory that was specified was not one of the following directories: /dumps, /dumps/iostats, /dumps/iotrace, /dumps/feature, /dumps/config, /dumps/elogs, /dumps/ec or /dumps/pl.

Explanation: An attempt was made to clear a file from, or copy a file to, a directory that is not valid.

User response: Ensure that the command accesses a valid directory.

CMMVC6204E The action failed as the resulting disk size would be less than, or equal to, zero.

Explanation: An attempt was made to shrink a disk, however the resulting size would have been less than or equal to zero.

User response: Not applicable

CMMVC6205E Metadata recovery can not use the provided MDisk id - invalid or destroyed.

Explanation: Metadata recovery cannot use the provided MDisk id, which is not valid or destroyed.

User response:

CMMVC6206E The software upgrade failed as a file containing the software for the specified MCP version was not found.

Explanation: There are two files required to successfully complete a software upgrade. One file contains the files that make up the base operating system, and the other file contains the 2145 software. This message appears if the OS version is incompatible with the 2145 software.

User response: Upload two compatible files, and resubmit the command.

CMMVC6207E The action failed because the virtual disk (VDisk) is part of a Remote Copy mapping.

Explanation: An action was performed against a volume that is part of a Remote Copy mapping.

User response: Remove the volume from the Remote

Copy mapping before resubmitting the command.

CMMVC6208E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.

Explanation: An action was performed against a volume that is part of a FlashCopy mapping.

User response: Remove the volume from the FlashCopy mapping before reissuing the command.

CMMVC6209E The FlashCopy mapping or consistency group could not be started in a reasonable time. The mapping or group is instead being prepared.

Explanation: The FlashCopy mapping or consistency group could not be started in a reasonable time. The mapping or group is instead being prepared.

User response: Resubmit the command.

CMMVC6210E The command has failed because a virtual medium error exists on the image mode VDisk or copy.

Explanation: When you submit this command, you cannot specify an image mode volume that has a virtual medium error on the volume or on any copy of the volume because the medium errors cannot be maintained on the ejected MDisk image copy.

User response: If an exact image copy is required, ensure that there is no virtual medium error on the image mode volume that you specify or on any of its copies, and resubmit the command.

If an exact copy is not required, you can use the `-force` option of the command, but all of the virtual medium errors will be lost.

CMMVC6211E The command failed as a migrate to image was in progress.

Explanation: An attempt was made to execute a command against a volume that was involved in a migrate to image operation.

User response: Wait for the migration to complete and reissue the command.

CMMVC6212E The command failed because data in the cache has not been committed to disk.

Explanation: The command failed because data in the cache has not been committed to disk.

User response:

CMMVC6213E You are trying to recover region data that was created by a code level different from the one you are currently running on the node.

Explanation: You are trying to recover region data that was created by a code level different from the one you are currently running on the node.

User response:

CMMVC6214E Failed to recreate the cluster you are trying to rebuild.

Explanation: Failed to recreate the cluster you are trying to rebuild.

User response:

CMMVC6215E The FlashCopy mapping was not created or modified because the consistency group already contains the maximum number of mappings.

Explanation: An attempt was made to create a FlashCopy mapping in, or move a FlashCopy mapping to, a consistency group that has the maximum number of FlashCopy mappings that it can contain.

User response: Create or move the FlashCopy mapping in another consistency group or remove an existing FlashCopy mapping from the desired group and then reissue the command.

CMMVC6216E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a Remote Copy mapping.

Explanation: The Remote Copy relationship was not created because the master or auxiliary volume is a member of a Remote Copy mapping.

User response: Select a different volume to make up the mapping.

CMMVC6217E The maximum number of hosts for the cluster is already configured.

Explanation: You must remove at least one host definition before you can resubmit the command.

User response: Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current host definitions are not required. Remove at least one host definition that is not required, and resubmit the command.

CMMVC6218E The maximum number of host/I/O group pairs for the cluster is already configured.

Explanation: You must remove at least one host I/O group pair definition before you can resubmit the command.

User response: Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current host I/O group pair definitions are not required. Remove at least one host I/O group pair definition that is not required, and resubmit the command.

CMMVC6219E The maximum number of WWPNS for the cluster is already configured.

Explanation: You must remove at least one WWPNS definition before you can resubmit the command.

User response: Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current WWPNS definitions are not required. Remove at least one WWPNS definition that is not required, and resubmit the command.

CMMVC6220E The maximum number of hosts for one or more IO groups is already configured.

Explanation: You must remove at least one host I/O group pair definition from the I/O group that you have specified before you can resubmit the command.

User response: Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current host I/O group pair definitions for the I/O group that you have specified are not required. Remove at least one host I/O group pair definition that is not required from the I/O group that you have specified, and resubmit the command.

CMMVC6221E The maximum number of WWPNS for one or more IO groups is already configured.

Explanation: You must remove at least one WWPNS definition from the I/O group that you have specified before you can resubmit the command.

User response: Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current

WWPN definitions for the I/O group that you have specified are not required. Remove at least one WWPNS definition that is not required from the I/O group that you have specified, and resubmit the command.

CMMVC6222E The maximum number of WWPNS for the host is already configured.

Explanation: You must remove at least one WWPNS definition for the host that you have specified before you can resubmit the command.

User response: Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current WWPNS definitions for the host that you have specified are not required. Remove at least one WWPNS definition that is not required for the host that you have specified, and resubmit the command.

CMMVC6223E The host does not belong to one or more of the IO groups specified or inferred.

Explanation: The host does not belong to one or more of the I/O groups specified or inferred.

User response: Specify a host I/O group combination that is currently defined, and resubmit the command.

CMMVC6224E The host already belongs to one or more of the IO groups specified.

Explanation: The host already belongs to one or more of the I/O groups specified.

User response: None.

CMMVC6225E An IO group cannot be removed from a host because of one or more associated VDisks.

Explanation: An I/O group cannot be removed from a host because of one or more associated volumes.

User response:

CMMVC6226E The action was not completed because the cluster has reached the maximum number of extents in MDisk groups.

Explanation: The cluster has reached the maximum number of extents in the storage pool; therefore, the action did not complete. You are attempting to use additional extents, for example by creating or expanding a volume. The action cannot be initiated because it would cause the maximum number of extents for a cluster to be exceeded.

User response: Free up extents by deleting other

volumes, and resubmit the command.

CMMVC6227I The package installed successfully.

Explanation: The package installed successfully.

User response: None.

CMMVC6228E The cluster was recovered and the CLI functionality is limited until the cause of the failure is determined and any corrective action taken. Contact IBM technical support for assistance.

Explanation: The cluster was recovered and the CLI functionality is limited.

User response: Contact IBM technical support for assistance.

CMMVC6229E The action failed as the SSH key has been revoked.

Explanation: The action failed as the SSH key has been revoked.

User response:

CMMVC6230E The action failed as the SSH key index (SSH_LABEL_ID) is invalid.

Explanation: The action failed as the SSH key index (SSH_LABEL_ID) is not valid.

User response:

CMMVC6231E The action failed as the audit table is full.

Explanation: The action failed as the audit table is full.

User response: Save the audit log to disk, and resubmit the command.

CMMVC6232E This operation cannot be performed because the cluster is currently aborting the previous software upgrade command.

Explanation: This operation cannot be performed because the cluster is currently cancelling the previous software upgrade command.

User response: Wait until the previous software upgrade command has cancelled successfully, and resubmit the command.

CMMVC6233E This operation cannot be performed because, either a software upgrade has not been started, or a software upgrade is in progress but is not in a state where it can be aborted.

Explanation: This operation cannot be performed because the software upgrade is making progress.

User response:

CMMVC6234E The upgrade cannot be aborted because at least one node has already committed to a new code level.

Explanation: The upgrade cannot be cancelled because at least one node has already committed to a new code level.

User response:

CMMVC6235E An invalid response has been entered. The command has not been executed. Input is case sensitive. Enter either yes or no.

Explanation: A response that is not valid has been entered. The command has not been executed.

User response: Enter either yes or no.

CMMVC6236E The command has not completed. A limited availability parameter has been entered without the required environment setting being set.

Explanation: The command has not completed. A limited availability parameter has been entered without the required environment setting being set.

User response:

CMMVC6237E The command failed as the remote cluster does not support global mirror.

Explanation: The command failed as the remote cluster does not support global mirror.

User response:

CMMVC6238E The copy type differs from other copies already in the consistency group.

Explanation: The copy type differs from other copies already in the consistency group.

User response: Ensure that the copy type of the mapping that you are attempting to add is the same copy type as the mappings in the consistency group to which you are attempting to add the mapping, and resubmit the command.

CMMVC6239E The FlashCopy mapping was not prepared because the mapping or consistency group is in the stopping state. The mapping or consistency group must first complete the stop operation and then be prepared

Explanation: You cannot prepare a FlashCopy mapping or consistency group when the FlashCopy mapping or consistency group is in the stopping state. If you want to prepare a FlashCopy mapping or consistency group, the FlashCopy mapping or consistency group must be in the Stopped or idle_or_copied state.

User response: Wait until the FlashCopy mapping or consistency group reaches the Stopped or idle_or_copied state and then resubmit the command.

CMMVC6240E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the stopping state.

Explanation: You cannot modify the consistency group of a FlashCopy mapping when the FlashCopy mapping is in the stopping state. If you want to modify the consistency group of a FlashCopy mapping, the FlashCopy mapping must be in the Stopped or idle_or_copied state.

User response: Wait until the FlashCopy mapping reaches the Stopped or idle_or_copied state and then resubmit the command.

CMMVC6241E The FlashCopy mapping was not deleted because the mapping or consistency group is in the stopping state. The mapping or consistency group must be stopped first.

Explanation: You cannot delete a FlashCopy mapping or consistency group when the FlashCopy mapping or consistency group is in the stopping state. If you want to delete a FlashCopy mapping or consistency group, the FlashCopy mapping or consistency group must be in the Stopped or idle_or_copied state.

User response: Wait until the FlashCopy mapping or consistency group reaches the Stopped or idle_or_copied state and then resubmit the command.

CMMVC6242E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopping state. The mapping or consistency group must first complete the stop operation and then be prepared.

Explanation: You cannot start a FlashCopy mapping or consistency group when the FlashCopy mapping or

consistency group is in the stopping state. If you want to start a FlashCopy mapping or consistency group, the FlashCopy mapping or consistency group must be in the Prepared state.

User response: Wait until the FlashCopy mapping or consistency group reaches the Stopped or idle_or_copied state and then prepare the FlashCopy mapping or consistency group before starting it.

CMMVC6243E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopping state.

Explanation: A Stop FlashCopy mapping or consistency group task has already been submitted and is still in progress. When the task has completed successfully, the FlashCopy mapping or consistency group state will change to Stopped.

User response: None.

CMMVC6244E The FlashCopy mapping was not created because the source virtual disk (VDisk) cannot be the target for a FlashCopy mapping.

Explanation: A volume cannot simultaneously be both the source of a FlashCopy mapping and the target of a FlashCopy mapping. The source volume that you have specified is currently defined as the target of a FlashCopy mapping.

User response: You have two options. One option is specify a different source volume and resubmit the command. The other option is delete the existing FlashCopy mapping that defines the source volume that you have specified as the target volume, and resubmit the command.

CMMVC6245E The FlashCopy mapping was not created because the source virtual disk (VDisk) is already in the maximum number of FlashCopy mappings.

Explanation: The number of FlashCopy mappings in which a volume can be defined as the source volume is limited. The source volume that you have specified cannot be defined to another FlashCopy mapping because it is already defined as the source volume to the maximum number of FlashCopy mappings.

User response: You have two options. One option is specify a different source volume and resubmit the command. The other option is delete one of the existing FlashCopy mappings that contains the source volume and resubmit the command.

CMMVC6246E The FlashCopy mapping was not created because the target virtual disk (VDisk) is already a source VDisk in a FlashCopy mapping.

Explanation: A volume cannot simultaneously be both the source of a FlashCopy mapping and the target of a FlashCopy mapping. The target volume that you have specified is currently defined as the source of a FlashCopy mapping.

User response: You have two options. One option is specify a different target volume and resubmit the command. The other option is delete all of the existing FlashCopy mappings that contain the target volume that you have specified and resubmit the command.

CMMVC6247E The FlashCopy mapping was not created because the target virtual disk (VDisk) is already a target VDisk in a FlashCopy mapping.

Explanation: A volume cannot simultaneously be the target of more than one FlashCopy mapping. The target volume that you have specified is currently defined as the target of another FlashCopy mapping.

User response: You have two options. One option is specify a different target volume and resubmit the command. The other option is delete the existing FlashCopy mapping that contains the target volume that you have specified and resubmit the command.

CMMVC6248E The command failed because the authorization table is full.

Explanation: The command failed because the authorization table is full.

User response:

CMMVC6249E The command failed because the authorization record was not found or is already set to the default role.

Explanation: The command failed because the authorization record was not found or is already set to the default role.

User response:

CMMVC6250E The command failed because the authorization record is not set to the default role. Use rmauth to set the default role.

Explanation: The command failed because the authorization record is not set to the default role.

User response: Use rmauth to set the default role.

CMMVC6251E The command failed because the specified role was not found.

Explanation: The command failed because the specified role was not found.

User response:

CMMVC6252E The command failed authorization because the session SSH key is invalid or was deleted.

Explanation: The command failed authorization because the session SSH key is not valid or was deleted.

User response:

CMMVC6253E The task has failed because the user's role is not authorized to submit the command.

Explanation: One example of a user role restriction is that a user that has a role of Monitor cannot create a volume.

User response: Either log in as a user that has a role that is authorized to submit the task or change the role of the user account that you are using to a role that is authorized to submit the task, and resubmit the task.

CMMVC6254E The command failed because the specified SSH key was not found.
NOTE This command must specify an admin key.

Explanation: The command failed because the specified SSH key was not found. This command must specify an admin key.

User response:

CMMVC6255E The command cannot set the authorization record to the default role. Use rmauth to set the default role.

Explanation: The command can not set the authorization record to the default role.

User response: Use rmauth to set the default role.

CMMVC6263E The command failed because the SSH key already exists or there is a duplicate SSH key.

Explanation: You have attempted to add an SSH key that already exists, and may have a different authorization level associated with it.

User response: Add a different SSH key if the existing SSH key of the same type does not have the authority level that you require.

CMMVC6269E Sendmail error EX_USAGE. A command or configuration line has been used incorrectly.

Explanation: The send email task has failed because a command or a configuration line has been used incorrectly.

User response: Ensure that the email settings are correct, and resubmit the task.

CMMVC6270E Sendmail error EX_DATAERR. Address is wrong, or the message is too large for the mailbox.

Explanation: The send email task has failed because the message sent is too large or a recipient address is incorrect.

User response: Ensure that all addresses are correct and that the message is not too large, and resubmit the task.

CMMVC6271E Sendmail error EX_NOINPUT. An input file (not a system file) did not exist or was not readable.

Explanation: The send email task has failed because a file is missing or cannot be read.

User response: Ensure that the email system is configured correctly. Ensure that access permissions have been specified correctly for all email configuration files, and resubmit the task.

CMMVC6272E Sendmail error EX_NOUSER. The sendmail command could not recognize a specified user.

Explanation: The send email task has failed because the user and domain combination that you specified does not exist.

User response: Specify a defined user and domain combination, and resubmit the task.

CMMVC6273E Sendmail error EX_NOHOST. The sendmail command could not recognize the specified host name.

Explanation: The send email task has failed because the host is not known to the email system.

User response: Ensure that you have configured the SMTP environment correctly. Ensure that you specify a defined host, and resubmit the task.

CMMVC6274E Sendmail error EX_UNAVAILABLE. A required system resource is not available.

Explanation: The send email task has failed because a required system resource is not available.

User response: Ensure that you have configured the SMTP environment correctly, and resubmit the task.

CMMVC6275E Sendmail error EX_SOFTWARE. An internal software error occurred (including bad arguments).

Explanation: The send email task has failed because an incorrect parameter or parameter value has been detected.

User response: Ensure that you have configured the SMTP environment correctly. Specify only supported parameters and parameter values, and resubmit the task.

CMMVC6276E Sendmail error EX_OSERR. A system resource error prevented the sending of an email.

Explanation: The send email task has failed because a system resource error has occurred.

User response: Ensure that you have configured the SMTP environment correctly, and resubmit the task.

CMMVC6277E Sendmail error EX_OSFILE. Failed to open a critical system file.

Explanation: The send email task has failed because a required system file cannot be opened.

User response: Ensure that the email system is configured correctly. Ensure that access permissions have been specified correctly for all email configuration files, and resubmit the task.

CMMVC6278E Sendmail error EX_CANTCREAT. An output file could not be written to by sendmail.

Explanation: The send email task has failed because the system cannot write to a required output file.

User response: Ensure that the email system is configured correctly. Ensure that access permissions have been specified correctly for all email configuration files, and resubmit the task.

CMMVC6279E Sendmail error EX_IOERR. A system I/O error occurred during a sendmail operation. This could be due to a disk failure.

Explanation: The send email task has failed because a write or read I/O operation has failed. This error might be caused by a disk device failure.

User response: Correct the root cause of the I/O failure, and resubmit the task.

CMMVC6280E Sendmail error EX_TEMPFAIL. The sendmail command could not create a connection to a remote system.

Explanation: The send email task has failed because the sendmail application cannot establish a connection to the remote system.

User response: Ensure that the network connection to the remote system is functioning correctly, and resubmit the task.

CMMVC6281E Sendmail error EX_PROTOCOL. The remote system returned something that was incorrect during a protocol exchange.

Explanation: The send email task has failed because an error in the protocol exchange has occurred.

User response: Ensure that the email system is configured correctly. Ensure that you have configured the SMTP environment correctly, and resubmit the task.

CMMVC6282E Sendmail error EX_NOPERM. The user does not have permission to perform the requested operation.

Explanation: The send email task has failed because the User ID does not have authorization to submit the task.

User response: Ensure that authorizations for your User ID in the email and SMTP configurations are correct, and resubmit the task.

CMMVC6283E Sendmail error EX_CONFIG. There is a fatal problem with the sendmail configuration.

Explanation: The send email task has failed because the sendmail configuration is not correct.

User response: Ensure that the email system is configured correctly. Ensure that you have configured the SMTP environment correctly, and resubmit the task.

CMMVC6284E An unknown error occurred. Please ensure your SMTP server is running.

Explanation: The send email task has failed because an unexpected error has occurred.

User response: Ensure that the SMTP server is running, and resubmit the task.

CMMVC6285E The email command timed out. Please check your email server settings as listed on the cluster.

Explanation: The send email task has failed because a command timeout has occurred.

User response: Ensure that your system settings match those recommended in the sendmail application documentation, and resubmit the task.

CMMVC6286E The email service has not been enabled.

Explanation: The send email task has failed because the email application is not enabled.

User response: Enable the email application, and resubmit the task.

CMMVC6287E The user specified does not exist.

Explanation: You must specify a User ID that exists.

User response: Ensure that the User ID that you specify is defined, and resubmit the task.

CMMVC6288E The FlashCopy mapping or consistency group could not be started because a source VDisk is the target of another FC Map that is keeping the VDisk inaccessible.

Explanation: You cannot start a FlashCopy mapping or consistency group when a source volume in the FlashCopy mapping or consistency group is the target volume of another FlashCopy mapping that is holding the volume as inaccessible. The task cannot be initiated because a source volume in the FlashCopy mapping or consistency group that you are attempting to start is the target of another FlashCopy mapping that is either prepared, preparing, stopped or stopping with a progress of less than 100%.

User response: Ensure that you have selected to start the correct FlashCopy mapping or consistency group. Ensure that none of the source volumes in the FlashCopy mapping or consistency group that you specify are target volumes in another FlashCopy mapping that is prepared, preparing, stopped or stopping with a progress of less than 100%. Resubmit the task.

CMMVC6289E The command failed because the virtual disk (VDisk) is pending synchronization.

Explanation: This error occurs when at least one of the volume copies is offline.

User response: Fix all of the errors that are associated with the volume copies, and resubmit the command.

CMMVC6290E The command failed because the virtual disk (VDisk) has image mode copies, is pending synchronization and -force has not been specified.

Explanation: This error occurs when at least one of

the volume copies is offline.

User response: Perform one of the following actions:

- Fix all of the errors that are associated with the volume copies, and resubmit the command.
- Resubmit the command and specify the `-force` parameter.

Note: When you specify the `-force` parameter with the command that caused this error, the image mode volume copy is no longer guaranteed to have the correct volume data.

CMMVC6291E The command failed because the virtual disk (VDisk) is pending synchronization and `-force` has not been specified.

Explanation: The command failed because the volume is pending synchronization and `-force` has not been specified.

User response: Perform one of the following actions:

- Fix all of the errors that are associated with the volume copies, and resubmit the command.
- Resubmit the command and specify the `-force` parameter.

Note: When you specify the `-force` parameter with the command that caused this error, the entire volume copy is resynchronized.

CMMVC6292E The command failed because a repair action is in progress for the virtual disk (VDisk).

Explanation: You cannot submit this command while a repair action is in progress for the volume.

User response: Use the `svcinfolrepairvdiskcopyprogress` command to view the repair progress. Wait for the volume repair process to complete. If you want the repair process to complete more quickly, increase the rate by submitting an `svctask chvdisk` command. Once the repair action has completed, resubmit the command that caused this error.

CMMVC6296E One or more managed disks (MDisks) have failed validation tests. The first failing MDisk ID is `MDISK_ID`.

Explanation: When you add a managed MDisk to a storage pool, the new MDisk is validated to ensure that adding it to the storage pool will not adversely impact the storage pool status. Either the current status of the MDisk has not allowed the validation to be performed, or the validation has failed. Note: You cannot add Read Only or faulty MDisks to a storage pool.

User response:

- If the MDisk identity has changed since it was last discovered, submit the command-line interface command `svctask detectmdisk`, which might correct the problem.
- Check switch zoning and logical unit presentation on the controller to ensure that the MDisk is physically and logically connected to all of the nodes in this cluster.
- Ensure that the controller settings are correct and that the MDisk logical unit is correctly configured.
- Ensure that the MDisk logical unit state is one that passes the validation. A Read Only or faulty MDisk fails the validation.
- View the cluster event log for more information about the failed validation.

CMMVC6297E One or more managed disks (MDisks) have timed out during validation tests. The first failing MDisk ID is `MDISK_ID`.

Explanation: When you add a managed MDisk to a storage pool, the new MDisk is validated to ensure that adding it to the storage pool will not adversely impact the storage pool status. The current status of the MDisk permits the validation to be initiated, but the allotted time for the validation process elapsed before the validation process had completed. Note: You cannot add Read Only or faulty MDisks to a storage pool.

User response:

- Ensure that the controller settings are correct and that the MDisk logical unit is correctly configured.
- Ensure that the MDisk logical unit state is one that passes the validation. A Read Only or faulty MDisk fails the validation.
- Check the fibre-channel fabric and storage controller for faults that might reduce the reliability of cluster communication with the MDisk.
- View the cluster event log for more information about the failed validation.

CMMVC6298E The command failed because a target VDisk has dependent FlashCopy mappings.

Explanation: The target volume of the FlashCopy mapping, or the target volume of at least one of the FlashCopy mappings in the consistency group, has other FlashCopy mappings that are dependent on the data on the target volume.

User response: Use the `lsvdiskdependentmaps` command and specify the target volume to determine which FlashCopy mappings are dependent on the target volume. Either wait for these mappings to reach the `idle_or_copied` state, or stop these mappings. Resubmit the command that produced this error.

CMMVC6299E The create failed because the source and target VDisks are members of FlashCopy mappings that have different grain sizes.

Explanation: All FlashCopy mappings that are in a tree of connected mappings must have the same grain size. The new FlashCopy mapping that you attempted to create would have linked two existing trees that have different grain sizes.

User response: You have three options. The first option is to resubmit the command and specify a different source or target volume. The second option is to delete all of the existing mappings that contain the source volume and resubmit the command. The third option is to delete all of the existing mappings that contain the target volume and resubmit the command.

CMMVC6300E The create failed because the source and target VDisks are members of FlashCopy mappings that belong to different I/O groups.

Explanation: All FlashCopy mappings in a tree of connected mappings must be in the same I/O group. The new FlashCopy mapping that you attempted to create would have linked two existing trees that are in different I/O groups.

User response: You have three options. The first option is to resubmit the command and specify a different source or target volume. The second option is to delete all of the existing mappings that contain the source volume and resubmit the command. The third option is to delete all of the existing mappings that contain the target volume and resubmit the command.

CMMVC6301E The create failed because the specified consistency group does not exist.

Explanation: The FlashCopy mapping was not created because the consistency group that you specified does not exist. You must create a consistency group before you can place a mapping in that group.

User response: Either create the FlashCopy consistency group that you specified and resubmit the command, or resubmit the command and specify an existing consistency group.

CMMVC6302E The create failed because the resulting tree of FlashCopy mappings would exceed the upper limit.

Explanation: Either the source volume or the target volume, or both, are already members of other FlashCopy mappings. The FlashCopy mapping was not created because the new FlashCopy mapping that you attempted to create would have linked two existing mapping trees into a single tree that exceeds the

maximum number of mappings that are supported for a single tree.

User response: You have two options. The first option is to resubmit the command and specify a different source or target volume. The second option is to delete a sufficient number of the existing FlashCopy mappings in which either the source or the target volume is a member so that the combined mapping tree does not exceed the maximum number of mappings that are supported for a single tree, and resubmit the command.

CMMVC6303E The create failed because the source and target VDisks are the same.

Explanation: A particular volume cannot be both the source and the target in a FlashCopy mapping. The FlashCopy mapping was not created because you have specified the same volume as both the source and the target.

User response: Resubmit the command and specify source and target volumes that are not identical.

CMMVC6304E The create failed because the source VDisk does not exist.

Explanation: You must specify an existing volume as the source of a FlashCopy mapping. The FlashCopy mapping was not created because the source volume that you specified does not exist.

User response: Either create the source volume that you specified and resubmit the command, or resubmit the command and specify an existing volume as the source.

CMMVC6305E The create failed because the target VDisk does not exist.

Explanation: You must specify an existing volume as the target of a FlashCopy mapping. The FlashCopy mapping was not created because the target volume that you specified does not exist.

User response: Either create the target volume that you specified and resubmit the command, or resubmit the command and specify an existing volume as the target.

CMMVC6306E The create failed because the source VDisk is the member of a FlashCopy mapping whose grain size is different to that specified.

Explanation: All FlashCopy mappings that are in a tree of connected mappings must have the same grain size. The FlashCopy mapping was not created because the source volume that you specified is either the source or the target volume of another FlashCopy mapping, and the grain size of the other mapping is different from the grain size that you specified for the

mapping that you attempted to create.

User response: You have two options. The first option is to delete all of the FlashCopy mappings that contain the source volume that you specified where the grain size of the FlashCopy mapping is different from the grain size that you specified, and resubmit the command. The second option is to resubmit the command and do not specify the grain size attribute.

CMMVC6307E The create failed because the target VDisk is the member of a FlashCopy mapping whose grain size is different to that specified.

Explanation: All FlashCopy mappings that are in a tree of connected mappings must have the same grain size. The FlashCopy mapping was not created because the target volume that you specified is either the source or the target volume of another FlashCopy mapping, and the grain size of the other mapping is different from the grain size that you specified for the mapping that you attempted to create.

User response: You have two options. The first option is to delete all of the FlashCopy mappings that contain the target volume that you specified where the grain size of the FlashCopy mapping is different from the grain size that you specified, and resubmit the command. The second option is to resubmit the command and do not specify the grain size attribute.

CMMVC6308E The create failed because the source VDisk is the member of a FlashCopy mapping whose IO group is different to that specified.

Explanation: All FlashCopy mappings in a tree of connected mappings must be in the same I/O group. The FlashCopy mapping was not created because the source volume that you specified is the source or target volume in another FlashCopy mapping and the I/O group of the other FlashCopy mapping is different from the I/O group that you specified.

User response: You have two options. The first option is to delete all of the FlashCopy mappings that contain the source volume that you specified where the FlashCopy mapping is in a different I/O group from the I/O group that you specified, and resubmit the command. The second option is to resubmit the command and do not specify the I/O group attribute. If you perform the second option, the default value of the I/O group attribute is used.

CMMVC6309E The create failed because the target VDisk is the member of a FlashCopy mapping whose IO group is different to that specified.

Explanation: All FlashCopy mappings in a tree of connected mappings must be in the same I/O group.

The FlashCopy mapping was not created because the target volume that you specified is the source or target volume in another FlashCopy mapping and the I/O group of the other FlashCopy mapping is different from the I/O group that you specified.

User response: You have two options. The first option is to delete all of the FlashCopy mappings that contain the target volume that you specified where the FlashCopy mapping is in a different I/O group from the I/O group that you specified, and resubmit the command. The second option is to resubmit the command and do not specify the I/O group attribute. If you perform the second option, the default value of the I/O group attribute is used.

CMMVC6310E The modify failed because the specified FlashCopy mapping does not exist.

Explanation: You cannot modify a FlashCopy mapping that does not exist. The modify command failed because the FlashCopy mapping that you specified does not exist.

User response: Resubmit the command and specify an existing FlashCopy mapping.

CMMVC6311E The command failed because the source VDisk is the target of a FlashCopy mapping that is in the specified consistency group.

Explanation: A particular volume cannot be both the source of one FlashCopy mapping and the target of another FlashCopy mapping in the same consistency group. The FlashCopy mapping was not created because the source volume of the FlashCopy mapping that you attempted to create is already the target volume of a FlashCopy mapping in the consistency group that you specified.

User response: Resubmit the command and specify a different consistency group.

CMMVC6312E The command failed because the target VDisk is the source of a FlashCopy mapping that is in the specified consistency group.

Explanation: A particular volume cannot be both the source of one FlashCopy mapping and the target of another FlashCopy mapping in the same consistency group. The FlashCopy mapping was not created because the target volume of the FlashCopy mapping that you attempted to create is already the source volume of a FlashCopy mapping in the consistency group that you specified.

User response: Resubmit the command and specify a different consistency group.

CMMVC6313E The command failed because the specified background copy rate is invalid.

Explanation: The command failed because the background copy rate that you specified is not a supported value.

User response: Either resubmit the command and specify a supported value for the background copy rate, or resubmit the command and do not specify the background copy rate attribute. If you do not specify the background copy rate attribute, the default background copy rate value is used.

CMMVC6314E The command failed because the specified cleaning rate is not valid.

Explanation: The command failed because the cleaning rate that you specified is not a supported value.

User response: Either resubmit the command and specify a supported value for the cleaning rate, or resubmit the command and do not specify the cleaning rate attribute. If you do not specify the cleaning rate attribute, the default cleaning rate value is used.

CMMVC6315E The command failed because the specified grain size is not valid.

Explanation: The command failed because the grain size that you specified is not a supported value.

User response: Either resubmit the command and specify a supported value for the grain size, or resubmit the command and do not specify the grain size attribute. If you do not specify the grain size attribute, the default grain size value is used.

CMMVC6319E The command has failed because a combination of IPv4 and IPv6 parameters were entered.

Explanation: The task accepts either IPv4 or IPv6 parameters. You cannot specify a combination of IPv4 and IPv6 parameters for this task.

User response: Specify only IPv4 or only IPv6 parameters, and resubmit the task.

CMMVC6320E The command has failed because the IPv4 address is not valid.

Explanation: The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

User response: Specify a valid IPv4 address, and resubmit the task.

CMMVC6321E The command has failed because the IPv4 subnet mask is not valid.

Explanation: The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

User response: Specify a valid IPv4 subnet mask, and resubmit the task.

CMMVC6322E The command has failed because the IPv4 gateway address is not valid.

Explanation: The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

User response: Specify a valid IPv4 gateway address, and resubmit the task.

CMMVC6323E The command has failed because the IPv6 address is not valid.

Explanation: Valid IPv6 address formats are:

- x:x:x:x:x:x:x
- x:x:x:x:x:d.d.d.d

where d is a decimal value from 0-255 of an IPv4 address and x is a hexadecimal value of an IPv6 address.

A special syntax is available to compress long strings of zero bits. The use of '::' indicates multiple groups of zeros. The ':' can appear only once in an address. The '::' can also be used to compress the leading or trailing zeros in an address.

- Example: 123.123.123.123
- Example: 1080:0:0:0:8:800:200C:417A, which can be compressed to 1080::8:800:200C:417A
- Example: 0:0:0:0:FFFF:129.144.52.38, which can be compressed to ::FFFF:129.144.52.38
- Example: 0:0:0:0:0:13.1.68.3, which can be compressed to ::13.1.68.3

User response: Specify a valid IPv6 address, and resubmit the task.

CMMVC6324E The command has failed because the IPv6 prefix is not valid.

Explanation: The value that you entered for an IPv6 address prefix is not a valid IPv6 address prefix.

User response: Specify a valid IPv6 address prefix, and resubmit the task.

CMMVC6325E The command has failed because the IPv6 gateway address is not valid.

Explanation: Valid IPv6 address formats are:

- x:x:x:x:x:x:x
- x:x:x:x:x:d.d.d.d

where d is a decimal value from 0-255 of an IPv4 address and x is a hexadecimal value of an IPv6 address.

A special syntax is available to compress long strings of zero bits. The use of '::' indicates multiple groups of zeros. The ':' can appear only once in an address. The '::' can also be used to compress the leading or trailing zeros in an address.

- Example: 123.123.123.123
- Example: 1080:0:0:0:8:800:200C:417A, which can be compressed to 1080::8:800:200C:417A
- Example: 0:0:0:0:FFFF:129.144.52.38, which can be compressed to ::FFFF:129.144.52.38
- Example: 0:0:0:0:0:13.1.68.3, which can be compressed to ::13.1.68.3

User response: Specify a valid IPv6 gateway address, and resubmit the task.

CMMVC6326E The command has failed because the IPv4 service state address is not valid.

Explanation: The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

User response: Specify a valid IPv4 service state address, and resubmit the task.

CMMVC6327E The command has failed because the IPv6 service state address is not valid.

Explanation: Valid IPv6 address formats are:

- x:x:x:x:x:x:x
- x:x:x:x:x:d.d.d.d

where d is a decimal value from 0-255 of an IPv4 address and x is a hexadecimal value of an IPv6 address.

A special syntax is available to compress long strings of zero bits. The use of '::' indicates multiple groups of zeros. The ':' can appear only once in an address. The '::' can also be used to compress the leading or trailing zeros in an address.

- Example: 123.123.123.123
- Example: 1080:0:0:0:8:800:200C:417A, which can be compressed to 1080::8:800:200C:417A
- Example: 0:0:0:0:FFFF:129.144.52.38, which can be compressed to ::FFFF:129.144.52.38
- Example: 0:0:0:0:0:13.1.68.3, which can be compressed to ::13.1.68.3

User response: Specify a valid IPv6 service state address, and resubmit the task.

CMMVC6328E The command has failed because the console address is not valid.

Explanation: The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

Valid IPv6 address formats are:

- x:x:x:x:x:x:x
- x:x:x:x:x:d.d.d.d

where d is a decimal value from 0-255 of an IPv4 address and x is a hexadecimal value of an IPv6 address.

A special syntax is available to compress long strings of zero bits. The use of '::' indicates multiple groups of zeros. The ':' can appear only once in an address. The '::' can also be used to compress the leading or trailing zeros in an address.

- Example: 123.123.123.123
- Example: 1080:0:0:0:8:800:200C:417A, which can be compressed to 1080::8:800:200C:417A
- Example: 0:0:0:0:FFFF:129.144.52.38, which can be compressed to ::FFFF:129.144.52.38
- Example: 0:0:0:0:0:13.1.68.3, which can be compressed to ::13.1.68.3

User response: Specify a valid console address, and resubmit the task.

CMMVC6329E The command has failed because the IP address is not valid.

Explanation: The valid IPv4 address format is d.d.d.d, where d is a decimal value from 0-255.

Valid IPv6 address formats are:

- x:x:x:x:x:x:x
- x:x:x:x:x:d.d.d.d

where d is a decimal value from 0-255 of an IPv4 address and x is a hexadecimal value of an IPv6 address.

A special syntax is available to compress long strings of zero bits. The use of '::' indicates multiple groups of zeros. The ':' can appear only once in an address. The '::' can also be used to compress the leading or trailing zeros in an address.

- Example: 123.123.123.123
- Example: 1080:0:0:0:8:800:200C:417A, which can be compressed to 1080::8:800:200C:417A
- Example: 0:0:0:0:FFFF:129.144.52.38, which can be compressed to ::FFFF:129.144.52.38
- Example: 0:0:0:0:0:13.1.68.3, which can be compressed to ::13.1.68.3

User response: Specify a valid IP address, and resubmit the task.

CMMVC6330E The command has failed because an IPv6 address was specified and the cluster does not have an IPv6 address.

Explanation: The cluster can only communicate with a server through an IPv6 address if an IPv6 cluster management IP address is configured.

User response: Either configure the cluster to have an IPv6 cluster management address or specify an IPv4 address, and resubmit the task.

Note: You do not need to remove the IPv4 address if you configure the cluster to have an IPv6 cluster management address.

CMMVC6331E The command has failed because an IPv4 address was specified and the cluster does not have an IPv4 address.

Explanation: The cluster can only communicate with a server through an IPv4 address if an IPv4 cluster management IP address is configured.

User response: Either configure the cluster to have an IPv4 cluster management address or specify an IPv6 address, and resubmit the task.

Note: You do not need to remove the IPv6 address if you configure the cluster to have an IPv4 cluster management address.

CMMVC6332E The command has failed because an IPv6 email server address was specified and the cluster does not have an IPv6 address.

Explanation: The cluster can only communicate with a server through an IPv6 address if an IPv6 cluster management IP address is configured.

User response: Either configure the cluster to have an IPv6 cluster management address or use an email server that has an IPv4 address, and resubmit the task.

Note: You do not need to remove the IPv4 address if you configure the cluster to have an IPv6 cluster management address.

CMMVC6333E The command has failed because an IPv4 email server address was specified and the cluster does not have an IPv4 address.

Explanation: The cluster can only communicate with a server through an IPv4 address if an IPv4 cluster management IP address is configured.

User response: Either configure the cluster to have an IPv4 cluster management address or use an email server that has an IPv6 address, and resubmit the task.

Note: You do not need to remove the IPv6 address if you configure the cluster to have an IPv4 cluster management address.

CMMVC6334E The command failed as the email port number supplied is invalid.

Explanation: The value that you entered for an email port number is not a valid email port number.

User response: Specify a valid email port number, and resubmit the task.

CMMVC6335E The command failed as the combination of parameters provided are either mutually incompatible or would leave the cluster without a functioning protocol stack.

Explanation: You have submitted a task with a combination of parameters and parameter values that is not supported or that does not provide the minimum amount of required information.

User response: Ensure that you specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6336E The virtual disk (VDisk) copy was not created because the grain size must be 32, 64, 128 or 256.

Explanation: You have supplied an incorrect value for the -grainsize parameter when you attempted to create a thin-provisioned volume copy.

User response: Specify a supported grain size, and resubmit the command.

CMMVC6337E The action failed because the warning size must be a multiple of 512 bytes.

Explanation: You are attempting to create a thin-provisioned volume copy but you have entered an incorrect value for the -warning parameter. The value can either be a percentage of the volume capacity or an absolute value that is a multiple of 512 bytes.

User response: Enter a supported warning value, and resubmit the command.

CMMVC6338E The action failed because the warning size can not be larger than the virtual size.

Explanation: You are attempting to create a thin-provisioned volume copy but you have entered an incorrect value for the -warning parameter. The warning value cannot be greater than the volume capacity.

User response: Enter a supported warning value, and resubmit the command.

CMMVC6339E The virtual disk (VDisk) copy was not created because the virtual size was not provided.

Explanation: You are attempting to create an image-mode thin-provisioned volume but you did not set the -size parameter.

User response: Resubmit the command using the -size parameter.

CMMVC6340E The action failed because the value supplied for real size is not a multiple of 512 bytes.

Explanation: You are attempting to create or resize a thin-provisioned volume copy but you have entered an incorrect value for the -rsize parameter. All sizes must be integer multiples of 512 bytes.

User response: Resubmit the command using a supported -rsize parameter value.

CMMVC6341E The action failed because the virtual disk (VDisk) copy is not space-efficient.

Explanation: You are attempting to run a command that is valid only for thin-provisioned volumes on a volume that is not thin-provisioned.

User response: Specify a thin-provisioned volume, and resubmit the command.

CMMVC6342E The virtual disk (VDisk) copy was not shrunk because its real size cannot be less than its used size.

Explanation: You are attempting to reduce the real size that is allocated to a thin-provisioned volume copy, but the command cannot be initiated because it would make the real size less than the size that is currently used.

User response: Determine the used size of the volume copy, and resubmit the command using a -rsize parameter value that is greater than or equal to the used size.

CMMVC6343E The virtual disk (VDisk) copy was not shrunk because its real size can not be negative.

Explanation: You are attempting to reduce the real size that is allocated to a thin-provisioned volume copy, but the command cannot be initiated because it would make the real size less than zero.

User response: Determine the real size of the volume copy, and resubmit the command using a supported -rsize parameter value.

CMMVC6344E The repair operation cannot start because the space-efficient virtual disk (VDisk) copy is already being repaired.

Explanation: You are attempting to repair a thin-provisioned volume copy, but the copy is already being repaired.

User response: Specify the correct volume and copy parameters, and resubmit the command.

CMMVC6345E The repair operation cannot start because the space-efficient virtual disk (VDisk) copy was created using -import but the cluster could not recognize its format.

Explanation: You are attempting to repair a thin-provisioned volume copy that is reporting corrupt metadata. The cluster cannot repair the volume copy because it was not recognized as a valid thin-provisioned volume when it was imported into this cluster. The most probable cause is that the wrong MDisk was used when the volume copy was imported.

User response: Delete the volume copy, and resubmit the import operation using the same MDisk that was exported from the original cluster.

CMMVC6346E The repair operation cannot start because the space-efficient virtual disk (VDisk) copy was created using -import with a real size that is too small.

Explanation: You are attempting to repair a thin-provisioned volume copy that is reporting corrupt metadata. The cluster cannot repair the volume copy because although it was recognized as a valid thin-provisioned volume when it was imported into this cluster, the real size allocated to the volume copy is too small. The most probable cause is that the incorrect value was supplied with -rsize parameter when the volume copy was imported.

User response: Delete the volume copy. Resubmit the import operation either using a larger value for -rsize, or supplying the -rsize parameter without a value to let the system choose a real size.

CMMVC6347E The specific upgrade package cannot be installed on this hardware level.

Explanation: The version of software that you are attempting to install does not support the hardware level of the configuration node.

User response: Check the release notes for the version of software that you want to install. Ensure that the version of software that you install supports the hardware level of all of the nodes in the cluster, and resubmit the task.

CMMVC6348E The command failed as there was not enough information provided to process successfully.

Explanation: You have submitted a task with a combination of parameters and parameter values that does not provide the minimum amount of required information.

User response: Ensure that you specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6349E The command was not initiated because the VDisk cache has been lost and you have not specified the `-force` option.

Explanation: You must specify the `-force` option when you move a volume from one I/O group to another and the volume has lost cache data.

User response: Resubmit the command and specify the `-force` option.

CMMVC6350E The command failed because there is insufficient mirror bitmap space.

Explanation: The command failed because there is insufficient free memory to allocate the bitmap needed for volume mirroring in the I/O group.

User response: Perform one of the following actions:

- Submit an `svctask chiogrp` command to increase the bitmap space.
- Remove volume mirrors from the I/O group.

Resubmit the command that caused this error.

CMMVC6351E The command failed because the virtual disk (VDisk) is not mirrored.

Explanation: Only mirrored volumes are supported for this command.

User response: Perform one of the following actions:

- Submit the appropriate command for a volume that is not mirrored.
- Submit an `svctask addvdiskcopy` command to add a copy to the volume, and resubmit the command that caused this error.

CMMVC6352E The command failed because the number of copies of this virtual disk (VDisk) would exceed the limit.

Explanation: You cannot exceed the limit on the number of copies that are supported for a volume.

User response: Submit an `svctask rmvdiskcopy` or `svctask splitvdiskcopy` command to decrease the

number of volume copies, and resubmit the command that caused this error.

CMMVC6353E The command failed because the copy specified does not exist.

Explanation: You must specify an existing copy for this command.

User response: Submit an `svcinfo lsvdiskcopy` command to show all of the available copies for this volume. Select a copy that exists, and then resubmit the command that caused this error.

CMMVC6354E The command failed because a copy is not synchronized.

Explanation: The copy that you specify for this command must be a synchronized copy.

User response: Use the `svcinfo lsvdisksyncprogress` command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting an `svctask chvdisk` command. When the copy is synchronized, resubmit the command that caused this error.

CMMVC6355E The command failed because an image mode copy is not synchronized and `-force` was not specified.

Explanation: When you specify an image mode copy for this command, the copy must be synchronized unless you also specify the `-force` parameter.

User response: Perform one of the following actions:

- Use the `svcinfo lsvdisksyncprogress` command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting an `svctask chvdisk` command. When the copy is synchronized, resubmit the command that caused this error.
- Resubmit the command and specify the `-force` parameter.

Note: When you specify the `-force` parameter with the command that caused this error, the image mode copy is no longer guaranteed to have the correct volume data.

CMMVC6356E The command failed because a copy is not synchronized and `-force` was not specified.

Explanation: When you specify a copy for this command, the copy must be synchronized unless you also specify the `-force` parameter.

User response: Perform one of the following actions:

- Use the **svcinfolsvdisksyncprogress** command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting an **svctask chvdisk** command. When the copy is synchronized, resubmit the command that caused this error.
- Resubmit the command and specify the **-force** parameter.

Note: When you specify the **-force** parameter with the command that caused this error, the entire volume copy is resynchronized.

CMMVC6357E The command failed because the copy specified is not synchronized and -force was not specified.

Explanation: When you specify a copy for this command, the copy must be synchronized unless you also specify the **-force** parameter.

User response: Perform one of the following actions:

- Use the **svcinfolsvdisksyncprogress** command to view the synchronization status. Wait for the copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting an **svctask chvdisk** command. When the copy is synchronized, resubmit the command that caused this error.
- Resubmit the command and specify the **-force** parameter.

Note: When you specify the **-force** parameter with the command that caused this error, the created volume is no longer guaranteed to have identical data to the original volume when the split is performed.

CMMVC6358E The command failed because the copy specified is the only synchronized copy.

Explanation: The command failed because the copy specified is the only synchronized copy.

User response: Use the **svcinfolsvdisksyncprogress** command to view the synchronization status. Wait for another copy to synchronize. If you want the synchronization process to complete more quickly, increase the rate by submitting an **svctask chvdisk** command. When the copy has synchronized, resubmit the command that caused this error.

CMMVC6359E The command failed because there are insufficient online synchronized copies.

Explanation: This error occurs when at least one of the volume copies is offline.

User response: Fix all of the errors that are associated

with the volume copies, and resubmit the command.

CMMVC6363E The command failed because the Logical Block Address (LBA) specified is invalid for this virtual disk (VDisk).

Explanation: You must specify a Logical Block Address (LBA) that is a valid address for this volume.

User response: Use the **svcinfolsvdisk** command to obtain the volume size, and resubmit the command that caused this error using a logical block address that is in range.

CMMVC6364E The command failed because the logical block address (LBA) requested is too large for the disk.

Explanation: You have specified an LBA in conjunction with a volume or MDisk, but the LBA is too large and does not exist on the disk.

User response: Check the size of the disk, and resubmit the command using an LBA that exists on the disk.

CMMVC6365E The command timed out.

Explanation: The command has not completed in a reasonable amount of time. Processing of the command required the software to wait for a set of MDisk reads or writes to complete, and the predefined reasonable wait time has been exceeded.

User response: Resolve any MDisk or fabric event log entries, and resubmit the command.

CMMVC6366E One or more nodes in the cluster has hardware that is not supported by the new software package.

Explanation: The version of software that you are attempting to install does not support the hardware in at least one node in the cluster.

User response: Check the release notes for the version of software that you want to install. Upgrade hardware so that all of the hardware in the cluster is supported by the new version of software, and resubmit the task.

CMMVC6367E A remote cluster is running software that is incompatible with the new software package.

Explanation: The version of software that you are attempting to install on the local cluster does not support the version of software that is installed on the remote cluster.

User response: Check the release notes for the version of software that you want to install. Perform one of the following actions:

CMMVC6368E • CMMVC6394E

- Upgrade the software on the remote cluster to a version that is supported by the version of software that you want to install on the local cluster before you upgrade the software on the local cluster.
- Delete the cluster partnership to stop all remote copy relationships between the clusters, and resubmit the task.

CMMVC6368E The new software package might be incompatible with the remote cluster.

Explanation: The software version compatibility between clusters cannot be checked because the remote cluster is not accessible.

User response: Perform one of the following actions:

- Ensure that the link to the remote cluster is functioning properly, and resubmit the task.
- Delete the cluster partnership to stop all remote copy relationships between the clusters, and resubmit the task.

CMMVC6369W The FlashCopy storage capacity that the cluster is using is approaching the FlashCopy storage capacity that is licensed.

Explanation: You are being warned that the FlashCopy storage capacity license might be exceeded soon.

User response: Upgrade the FlashCopy storage capacity license to prevent recurrence of this warning message.

CMMVC6370W The Remote Copy storage capacity that the cluster is using is approaching the Remote Copy storage capacity that is licensed.

Explanation: You are being warned that the Remote Copy storage capacity license might be exceeded soon.

User response: Upgrade the Remote Copy storage capacity license to prevent recurrence of this warning message.

CMMVC6372W The virtualized storage capacity that the cluster is using is approaching the virtualized storage capacity that is licensed.

Explanation: You are being warned that the virtualized storage capacity license might be exceeded soon.

User response: Upgrade the virtualized storage capacity license to prevent recurrence of this warning message.

CMMVC6373W The virtualized storage capacity that the cluster is using exceeds the virtualized storage capacity that is licensed.

Explanation: You are being warned that the virtualized storage capacity license has been exceeded.

User response: Upgrade the virtualized storage capacity license to prevent recurrence of this warning message.

CMMVC6374W The FlashCopy storage capacity that the cluster is using exceeds the FlashCopy storage capacity that is licensed.

Explanation: You are being warned that the FlashCopy storage capacity license has been exceeded.

User response: Upgrade the FlashCopy storage capacity license to prevent recurrence of this warning message.

CMMVC6375W The Remote Copy storage capacity that the cluster is using exceeds the Remote Copy storage capacity that is licensed.

Explanation: You are being warned that the Remote Copy storage capacity license has been exceeded.

User response: Upgrade the Remote Copy storage capacity license to prevent recurrence of this warning message.

CMMVC6394E The command failed because an attempt to make the virtual disk cache empty took too long.

Explanation: The failed command must empty the volume cache before attempting the requested action to ensure that data is preserved. The empty volume cache subtask has taken too long, and therefore the command that you have submitted was not initiated so that other configuration activity can occur.

The system continues attempting to empty the volume cache.

The storage associated with the volume is probably overloaded.

User response: Wait a few minutes to allow the volume cache to empty. Resubmit the command.

Alternatively, you can use the `-force` parameter, if the command supports the `-force` parameter, to bypass the empty volume cache subtask. However, specifying the `-force` parameter will discard cache data for the volume. Only use the `-force` flag with this command if you do not intend to use the existing contents of the volume.

In addition to the above actions, investigate the performance of the network storage devices associated with this volume. The performance of host applications using these devices might be degraded.

Remedial action to resolve a performance problem enables host application performance to return to optimal conditions, and prevents this error message from recurring when you resubmit the command that caused this error.

CMMVC6399E The command failed because there is not enough memory available for reservation.

Explanation: At least one node in the cluster cannot reserve the required amount of memory. This might be caused by pinned data in the cache.

User response: Check for events in the event log. Follow the directed maintenance procedures to resolve the problem.

CMMVC6400E The command failed because a specified managed disk (MDisk) is already in use.

Explanation: You cannot specify an MDisk for this command if it is already in a storage pool or is being used as an image mode volume.

User response: Specify an MDisk that is not being used as an image mode volume and is not in a storage pool, and resubmit the command.

CMMVC6401E The command failed because one or more of the specified managed disks (MDisks) that you have specified are not in the required managed disk group.

Explanation: The command requires that all of the MDisks that you specify must be in the same storage pool.

User response: Ensure that all of the MDisks that you specify are in the same storage pool, and resubmit the command.

CMMVC6402E The command failed because the managed disk (MDisk) is not in the required managed disk group.

Explanation: All of the MDisks that you specify must be in the required storage pool. At least one of the source MDisks that you have specified in the command is not in the required storage pool.

User response: Ensure that all of the MDisks that you specify are in the storage pool that you specify, and resubmit the command.

CMMVC6403E The command failed because the target managed disk (MDisk) is not in the required managed disk group.

Explanation: All of the MDisks that you specify must be in the required storage pool. At least one of the target MDisks that you have specified in the command is not in the required storage pool.

User response: Ensure that all of the MDisks that you specify are in the storage pool that you specify, and resubmit the command.

CMMVC6404E The command failed because the source and target managed disk groups must be different.

Explanation: The source and target storage pools that you specify for a cross storage pool migration must be different.

User response: Ensure that the source and target storage pools that you specify for a cross storage pool migration are different, and resubmit the command.

CMMVC6405E The command failed because the target copy was not specified.

Explanation: A target copy must be specified when you use migrations on a volume and more than one volume copy exists.

User response: Specify the target copy, and resubmit the command.

CMMVC6406E The command failed because the specified managed disk group does not exist.

Explanation: At least one of the storage pools that you have specified in the parameter list does not exist.

User response: Ensure that each of the storage pools that you specify exists, and resubmit the command.

CMMVC6407E The command failed because the managed disk group is invalid.

Explanation: At least one storage pool ID is above the maximum value that is available for the system.

User response: Ensure that each storage pool ID that you specify in the parameter list exists, and resubmit the command.

CMMVC6408E The command failed because too few managed disk groups were specified.

Explanation: You must specify the number of storage pools that is consistent with the other parameters and parameter values that you specify with the command.

User response: Refer to the command documentation

for valid combinations of parameters and parameter values. Use a valid combination of parameters and values, and resubmit the command.

CMMVC6409E The command failed because too many managed disk groups were specified.

Explanation: You must specify the number of storage pools that is consistent with the other parameters and parameter values that you specify with the command.

User response: Refer to the command documentation for valid combinations of parameters and parameter values. Use a valid combination of parameters and values, and resubmit the command.

CMMVC6410E The command failed because too few managed disks (MDisks) were specified.

Explanation: You must specify the number of MDisks that is consistent with the other parameters and parameter values that you specify with the command.

User response: Refer to the command documentation for valid combinations of parameters and parameter values. Use a valid combination of parameters and values, and resubmit the command.

CMMVC6411E The command failed because too many managed disks (MDisks) were specified.

Explanation: You must specify the number of MDisks that is consistent with the other parameters and parameter values that you specify with the command.

User response: Refer to the command documentation for valid combinations of parameters and parameter values. Use a valid combination of parameters and values, and resubmit the command.

CMMVC6412E The command failed because the managed disk group extent size is above maximum permitted size.

Explanation: You cannot specify a storage pools extent size that is larger the maximum size.

User response: Specify a storage pool extent size that is less than or equal to the maximum size, and resubmit the command.

CMMVC6413E The command failed because the managed disk (MDisk) is invalid.

Explanation: At least one MDisk ID is above the maximum value that is available for the system.

User response: Ensure that each MDisk ID that you specify in the parameter list exists, and resubmit the command.

CMMVC6414E The command failed because the managed disk (MDisk) is currently being migrated.

Explanation: When you submit this command, you cannot specify an MDisk that is being migrated.

User response: Either wait until the migration has completed for the MDisk that you specify, or specify a different MDisk, and resubmit the command.

CMMVC6415E The command failed because the managed disk group warning threshold is too low.

Explanation: You must specify a storage pool warning threshold that is equal to or greater than the minimum size.

User response: Specify a storage pool warning threshold that is equal to or greater than the minimum size, and resubmit the command.

CMMVC6416E The command failed because the managed disk group warning threshold is too high.

Explanation: You must specify a storage pool warning threshold size that is equal to or less than the size of the storage pool when all of the MDisks have been added, or you must specify a storage pool warning percentage that is equal to or less than the maximum warning threshold percentage.

User response: Specify valid values for the storage pool warning threshold size or percentage, and resubmit the command.

CMMVC6417E The command failed because the managed disk group warning threshold is invalid.

Explanation: To specify the warning threshold there must be at least one managed MDisk in the storage pool.

User response: Ensure that there is at least one MDisk defined for the storage pool or remove the warning threshold, and resubmit the command.

CMMVC6418E The command failed because the virtual disk (VDisk) is in the process of being resized.

Explanation: When you submit this command, you cannot specify a volume that is being resized.

User response: Wait for the resize volume operation to complete. If you still want to submit this command after the operation has completed, resubmit the command.

CMMVC6419E The command failed because one or more of the specified managed disks (MDisks) are in the process of being deleted.

Explanation: When you submit this command, you cannot specify an MDisk that is being deleted with the -force option.

User response: Wait for the delete MDisk operation to complete. Do not include any MDisks that have been deleted in the list of MDisks that you specify, and resubmit the command.

CMMVC6423E The Send Inventory email operation failed because email is not started.

Explanation: The send inventory email functionality has been enabled but the email service has not been started.

User response: Disable the send inventory email functionality or start the email service.

CMMVC6424E The Send Inventory email operation failed because there are no inventory email users.

Explanation: The send inventory functionality has been enabled but no email users with the ability to receive inventory emails have been created.

User response: Either turn off the send inventory email functionality or create an email user account that is capable of receiving inventory emails. Refer to the documentation for the mke-mailuser command for help on creating email users.

CMMVC6425E The action failed because the maximum number of objects has been reached.

Explanation: The action failed because the maximum number of objects has been reached.

User response: Not applicable.

CMMVC6426E The command failed because a specified managed disk (MDisk) is already in use.

Explanation: You cannot specify an MDisk that is already configured as an image mode volume.

User response: Specify an unmanaged disk, and resubmit the task.

CMMVC6427E The command failed because one or more of the specified managed disks (MDisks) are not in the required managed disk group.

Explanation: The create volume task requires that all

of the MDisks that you specify must be in the same storage pool.

User response: Ensure that all of the MDisks that you specify are in the same storage pool, and resubmit the task.

CMMVC6428E The command failed because the source managed disk (MDisk) is not in the required managed disk group.

Explanation: The task requires that all of the source MDisks that you specify must be in the same storage pool.

User response: Ensure that all of the source MDisks that you specify are in the same storage pool, and resubmit the task.

CMMVC6429E The command failed because the target managed disk (MDisk) is not in the required managed disk group.

Explanation: The task requires that all of the target MDisks that you specify must be in the same storage pool.

User response: Ensure that all of the target MDisks that you specify are in the same storage pool, and resubmit the task.

CMMVC6430E The command failed because the target and source managed disk groups must be different.

Explanation: The cross storage pool migration task does not support specifying the same storage pool to be both the source and target storage pool.

User response: Specify a source storage pool and a target storage pool that are not identical, and resubmit the task.

CMMVC6431E The command failed because the target copy was not specified.

Explanation: When you use migrations on a volume and there is more than one copy, you must specify which copy to use as the target copy.

User response: Specify the target copy, and resubmit the task.

CMMVC6432E The command failed because the specified managed disk group does not exist.

Explanation: All of the storage pools that you specify must already exist.

User response: Ensure that all of the storage pools that you specify already exist, and resubmit the task.

CMMVC6433E The command failed because the managed disk group is invalid.

Explanation: All of the storage pool IDs that you specify must have a value that is less than or equal to the maximum supported storage pool ID value.

User response: Ensure that all storage pools have supported ID values. Ensure that all of the storage pools that you specify already exist, and resubmit the task.

CMMVC6434E The command failed because too few managed disk groups were specified.

Explanation: The combination of parameters and parameter values that you have specified is not supported. The task requires that you specify more storage pools than the number that you have specified.

User response: Specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6435E The command failed because too many managed disk groups were specified.

Explanation: The combination of parameters and parameter values that you have specified is not supported. The task requires that you specify fewer storage pools than the number that you have specified.

User response: Specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6436E The command failed because too few managed disks (MDisks) were specified.

Explanation: The combination of parameters and parameter values that you have specified is not supported. The task requires that you specify more MDisks than the number that you have specified.

User response: Specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6437E The command failed because too many managed disks (MDisks) were specified.

Explanation: The combination of parameters and parameter values that you have specified is not supported. The task requires that you specify fewer MDisks than the number that you have specified.

User response: Specify a supported combination of parameters and parameter values, and resubmit the task.

CMMVC6438E The command failed because the managed disk group extent size is above maximum permitted size.

Explanation: The storage pool extent size that you have specified is greater than the supported maximum value.

User response: Specify a supported storage pool extent size, and resubmit the task.

CMMVC6439E The command failed because the managed disk (MDisk) is invalid.

Explanation: Each MDisk ID must have a value that is less than or equal to the maximum supported MDisk ID value.

User response: Ensure that all of the MDisks have supported ID values. Ensure that all of the MDisks that you specify already exist, and resubmit the task.

CMMVC6440E The command failed because the managed disk (MDisk) is currently being migrated.

Explanation: When you submit this task, you cannot specify an MDisk that is being migrated.

User response: Ensure that the MDisk that you specify is not migrating, and resubmit the task. If you want to specify the same MDisk and resubmit the task, ensure that the migration for that MDisk has completed before you resubmit the task.

CMMVC6441E The command failed because the managed disk group warning threshold is too low.

Explanation: The value that you have specified for the storage pool warning threshold is less than the minimum supported value.

User response: Specify a supported value for the storage pool warning threshold, and resubmit the task.

CMMVC6442E The command failed because the managed disk group warning threshold is too high.

Explanation: Either the value for the storage pool warning percentage is greater than the maximum supported value, or the storage pool warning disk size is greater than the storage pool capacity.

User response: Specify supported values for storage pool warning percentage and disk size, and resubmit the task.

CMMVC6443E The command failed because the managed disk group warning threshold is invalid.

Explanation: If you submit this command and specify a storage pool warning threshold percentage, you must specify a storage pool that contains at least one MDisk and you must specify a supported value for the storage pool warning threshold percentage.

User response: Either do not specify a storage pool warning threshold percentage, or specify a supported value for the storage pool warning threshold percentage and specify a storage pool that contains at least one MDisk, and resubmit the task.

CMMVC6444E The command failed because the virtual disk (VDisk) is in the process of being resized.

Explanation: You cannot specify a volume that is being resized when you submit this task.

User response: Wait for the resize volume task to complete. You can specify the same volume and resubmit this task only after the resize volume task that is in progress has completed.

CMMVC6445E The command failed because one or more of the specified managed disks (MDisks) are in the process of being deleted.

Explanation: You cannot specify an MDisk that is being force deleted.

User response: Wait until all force delete MDisk tasks have completed. Ensure that all of the MDisks that you specify still exist, and resubmit the task.

CMMVC6446E The command failed because the managed disk groups have different extent sizes.

Explanation: This task requires that the extent size of the source storage pool and the extent size of the target storage pool must be identical.

User response: If you want to resubmit this command, ensure that the source and target storage pools have the same extent size. If you want to move a volume to a storage pool that has a different extent size, you must use the procedure that is documented in the technical notes.

CMMVC6447E The command failed because the virtual disk (VDisk) is currently being migrated.

Explanation: You cannot specify a volume that is being migrated.

User response: Either wait until the volume migration

process has completed and resubmit the task, or specify a volume that is not being migrated and resubmit the task.

CMMVC6448E Deleting this node will cause data loss for resources associated with the I/O group of this node.

Explanation: This node contains resources which are vital to the I/O group and unavailable elsewhere. Removing this node will cause a loss of customer data.

It is recommended that this node not be removed unless the customer data supported by it is of no consequence.

User response: The -force option must be used to remove this node.

CMMVC6449E The operation was not performed because the partnership owns Global or Metro Mirror relationships or consistency groups.

Explanation: The cluster partnership cannot be removed while there are Global or Metro Mirror relationships or consistency groups that are configured in the local cluster and that are associated with the remote cluster of the partnership.

User response: Identify all of the Global or Metro Mirror relationships or consistency groups in the local cluster that are configured between this cluster and the remote cluster of the partnership. Remove all of the relationships and groups that you have identified, and resubmit the task.

Note: Do not remove relationships or groups that are associated with a different cluster, and do not remove relationships or groups that are contained entirely within the local cluster.

CMMVC6450W A FlashCopy mapping was created but physical_flash is not enabled.

Explanation: The create FlashCopy mapping task has succeeded. However, physical_flash should be enabled when you create a FlashCopy mapping in the physical disk license scheme.

User response: Ensure that you have the appropriate virtualization license for the cluster configuration that you want to enable. Ensure that the license settings for this cluster match the license.

Delete the FlashCopy mapping or enable physical_flash.

CMMVC6451W A Global Mirror or Metro Mirror relationship was created but `physical_remote` is not enabled.

Explanation: The create Global Mirror or Metro Mirror relationship task has succeeded. However, `physical_remote` should be enabled when you create a Global Mirror or Metro Mirror relationship and the cluster uses the physical disk license scheme.

User response: Ensure that you have the appropriate virtualization license for the cluster configuration that you want to enable. Ensure that the license settings for this cluster match the license.

Delete the Global Mirror or Metro Mirror relationship or enable `physical_remote`.

CMMVC6452W You are using the physical disk license scheme but the values for `physical_flash` and `physical_remote` are not set.

Explanation: The task has succeeded. However, you should enable `physical_flash` before you create a FlashCopy mapping and you should enable `physical_remote` before you create a Global Mirror or Metro Mirror mapping.

User response: Enable `physical_flash` before you create a FlashCopy mapping. Enable `physical_remote` before you create a Global Mirror or Metro Mirror relationship.

CMMVC6453W You have disabled the physical disk license scheme but the capacity license scheme is not set.

Explanation: The task has succeeded. However, you should configure a license scheme before you create a FlashCopy, Global Mirror or Metro Mirror relationship. You can configure a physical disk license scheme or a capacity license scheme, but not both.

User response: If you do not have a virtualization feature license that is valid for this cluster, contact your IBM sales representative and obtain a license. Ensure that the license settings for this cluster match the license that you have for this cluster.

CMMVC6454E The command failed because the physical disk license scheme is not enabled.

Explanation: You can only enable `physical_flash` or `physical_remote` when the physical disk license scheme is enabled.

User response: Ensure that you have the appropriate virtualization license for the cluster configuration that you want to enable. Ensure that the license settings for this cluster match the license. Resubmit the task if it supported by the license.

CMMVC6455E The command failed because a capacity license scheme parameter was specified but the physical disk license scheme is enabled.

Explanation: You cannot enable the capacity license scheme or specify a capacity license scheme parameter while the cluster is using the physical disk license scheme.

User response: Ensure that you have the appropriate virtualization license for the cluster configuration that you want to enable. Ensure that the license settings for this cluster match the license. Resubmit the task if it supported by the license.

CMMVC6456E The command failed because a physical disk license scheme parameter was specified but the capacity license scheme is enabled.

Explanation: You cannot enable the physical disk license scheme or specify a physical disk license scheme parameter while the cluster is using the capacity license scheme.

User response: Ensure that you have the appropriate virtualization license for the cluster configuration that you want to enable. Ensure that the license settings for this cluster match the license. Resubmit the task if it supported by the license.

CMMVC6457E One or more quorum disks are on the specified controller.

Explanation: You cannot disable the setting that allows a controller to support a quorum disk while a quorum disk is configured on the controller.

User response: Move all quorum disks from the controller to a different storage system using the `svctask setquorum` command, and resubmit this task.

CMMVC6458E The specified controller cannot support quorum disks.

Explanation: The controller type of the controller that you specified does not support quorum disks.

User response: Specify a controller that has a controller type that supports quorum disks, and resubmit the task.

CMMVC6459E The `mkcrrelationship` command failed because the same VDisk was specified as the master and auxiliary VDisk.

Explanation: A relationship cannot be created from a volume to itself. The `mkcrrelationship` command requires that you specify two different volumes for the master and auxiliary positions. These can be two

volumes in the local cluster, or a volume in each of two different clusters.

User response: Specify a master volume and an auxiliary volume that are not identical to each other, and resubmit the task.

CMMVC6460E The command failed because the migration source is offline.

Explanation: The source of the migration is offline. The offline source is either an image mode MDisk or the entire storage pool.

User response:

- If you submitted the `rmmdisk` command and specified a regular MDisk, determine the storage pool to which the source MDisk is defined, and follow the procedure for bringing the storage pool online. There will be an entry in the event log for the corresponding storage pool.
- If you submitted the `rmmdisk` command and specified an image mode MDisk, determine the source MDisk and follow the procedure for bringing the image mode MDisk online. There will be an entry in the event log for the corresponding MDisk.
- If you submitted a command to migrate a copy of an image mode volume, determine the corresponding source MDisk and follow the procedure for diagnosing problems with MDisks. There will be an entry in the event log for the corresponding MDisk.
- If you submitted any other command to migrate a volume copy, determine the storage pool to which the volume is defined, and follow the procedure for bringing the storage pool online. There will be an entry in the event log for the corresponding storage pool.

CMMVC6461E The command failed because starting the migration will result in VDIsks going offline in the source managed disk group.

Explanation: A migration from an image mode volume will use the source storage pool and the source storage pool assumes the combined state of the image mode MDisk and the storage pool. If the online or offline states of the image mode MDisk and the storage pool are different on different nodes, the source volume might go offline or all of the volumes in the source storage pool might go offline.

User response: For each node, note the online or offline states of the source volume and the source storage pool. If one entity is online and the other is offline, bring online whichever is offline. Taking the online entity offline is not recommended because other volumes might go offline.

CMMVC6462E The command failed because starting the migration will result in VDIsks going offline because the target managed disk group is offline.

Explanation: The migration process assigns the volume an online or offline state based on the states of the source and target storage pools. In this case, based on the offline state of the target storage pool the volume that is currently online would have been taken offline. The command cannot be initiated because this action is not supported. There will be an entry in the event log for the corresponding storage pool.

User response: For each node, note the online or offline state of the source and target storage pools. For each node, if one of these two storage pools is online and the other is offline, bring online whichever storage pool is offline. Taking the online storage pool offline is not recommended because other volumes might go offline.

CMMVC6463E The command failed because Starting the migration will result in VDIsks going offline because a target MDisk is offline.

Explanation: The volume is currently online. The migration process assigns the volume an online or offline state based on the states of the source and target MDisks. In this case, based on the offline state of the target MDisk, the volume would have been taken offline. The task cannot be initiated because this action is not supported.

User response: Bring the target MDisk online by following the recommended procedure for bringing an MDisk online, and resubmit the command.

CMMVC6464E The Create FlashCopy mapping task cannot be initiated because the size of the source VDisk is being changed by a previously submitted task.

Explanation: You cannot submit this task while the Change volume size task is in progress.

User response: Wait until the Change volume size task completes, and then resubmit the task.

CMMVC6465E The Create FlashCopy mapping task cannot be initiated because the size of the target VDisk is being changed by a previously submitted task.

Explanation: You cannot submit this task while the Change volume size task is in progress.

User response: Wait until the Change volume size task completes, and then resubmit the task.

CMMVC6466E The Create FlashCopy mapping task cannot be initiated because an identical map already exists.

Explanation: A map between the source and target volumes that you have specified is defined. You cannot define a map that is exactly the same as a map that is already defined.

User response: Specify a unique map when you submit this task.

CMMVC6467E The Create FlashCopy mapping task cannot be initiated because a FlashCopy map with the same target VDisk already exists in the consistency group.

Explanation: You cannot create more than one FlashCopy map with the same target volume in the same consistency group.

User response: Specify a target volume for the FlashCopy map that is unique to the consistency group when you submit this task.

CMMVC6468E The Start or Prepare FlashCopy mapping task cannot be initiated because the target volume is the source of a different FlashCopy map that is being restored.

Explanation: You cannot start or prepare a map while the target of the map is the source volume of another FlashCopy mapping that is being restored.

User response: Ensure that the target volume in the map that you are attempting to start or prepare is not the source volume of another FlashCopy mapping that is being restored when you submit the task. You could stop the associated map that is being restored, or you could wait for the map that is being restored to reach the Idle_or_Copied state.

CMMVC6469E The Split stop FlashCopy map task cannot be initiated because the mapping is either being restored or is not in the copy complete state.

Explanation: You cannot split stop a FlashCopy map while it is being restored or is not in the copy complete state.

User response: Ensure that the map is not being restored and is in the copy complete state when you submit this task.

CMMVC6470E The Start or Prepare FlashCopy mapping task cannot be initiated because the target VDisk is being used by a different FlashCopy map.

Explanation: You cannot start or prepare a map while the target of the map is also the target volume of

another map that is in one of the following states: copying, stopping, suspended, prepared or preparing.

User response: Ensure that the target volume in the map that you are attempting to start or prepare is not the target volume of another FlashCopy mapping that is in one of the unsupported states when you submit this task.

CMMVC6471E The Create cluster partnership task cannot be initiated because a cluster in the existing partnership has a downlevel software version that does not support this configuration.

Explanation: One scenario in which this error occurs is when a cluster at version 5.1.0 or later is partnered to a cluster at version 4.3.1 or earlier and you attempt to create another partnership to a cluster at version 5.1.0 to implement multiple cluster mirroring. Software version 4.3.1 does not support multiple cluster mirroring, so adding a partnership to a third cluster is not supported while at least one cluster in the current partnership is at version 4.3.1 or earlier.

User response: Either upgrade the downlevel cluster software version to a version that supports this task or remove the partnership to the cluster that has the downlevel software version, and resubmit the task.

CMMVC6472E The Create cluster partnership task cannot be initiated because the remote cluster with which you are attempting to create a partnership has a downlevel software version that does not support this configuration.

Explanation: The software versions of the clusters in the existing partnership do not support a partnership with a cluster that has the software version of the remote cluster with which you are attempting to create a partnership. If a cluster at version 5.1.0 or later is already in a partnership with another cluster at version 5.1.0 or later, you can only add a partnership to a cluster at version 5.1.0 or later, and cannot add a partnership to a cluster at version 4.3.1 or earlier. If a cluster at version 5.1.0 or later is already in a partnership with another cluster at version 4.3.1 or earlier, you cannot add another partnership while the partnership with the cluster at version 4.3.1 exists. If a cluster is not in a partnership, you can create a partnership between it and a cluster at any version. One scenario in which this error occurs is when you attempt to add a partnership with a remote cluster at version 4.3.1 or earlier to a cluster at software version 5.1.0 or later that is already in partnership with another cluster at software version 5.1.0 or later.

User response: Either upgrade the downlevel cluster software version to a version that supports this task or remove all existing partnerships from the cluster to which you want to partner the cluster that has the

downlevel software version, and resubmit the task.

CMMVC6473E The partnership task cannot be initiated because the supported maximum number of accessible remote clusters would be exceeded.

Explanation: With multiple cluster mirroring, you can build a configuration of a chain of clusters. There is a limit to the number of clusters that you can configure in the chain. The task would have resulted in exceeding the supported maximum number of clusters in a chain.

User response: Ensure that the resulting configuration is supported when you submit this task.

CMMVC6474E The Create partnership task cannot be initiated because there is a Global Mirror or Metro Mirror relationship or consistency group that has a deleted partnership.

Explanation: You must resolve the unpartnered objects error that is related to the deleted partnership with a Global Mirror or Metro Mirror relationship or consistency group before you can create a partnership from the local cluster to more than one other cluster.

User response: Resolve the unpartnered objects error, and resubmit the task. To resolve the error, either delete the unpartnered Global Mirror or Metro Mirror relationship or consistency group from the deleted partnership, or create a partnership for the unpartnered objects.

CMMVC6475E The Add relationship to group task cannot be initiated because the master cluster of the relationship that you are attempting to add to the group is the auxiliary cluster of the group, and the auxiliary cluster of the relationship that you are attempting to add to the group is the master cluster of the group.

Explanation: All of the relationships within a group must have the same master cluster as the group and must have the same auxiliary cluster as the group. The determination as to which cluster is assigned as the master cluster when you create a relationship or consistency group is based on the cluster from which you submit the task.

User response: Perform one of the following three options:

- Delete the group and create the group so that the master cluster of the group is identical to the master cluster of the relationship and the auxiliary cluster of the group is identical to the auxiliary cluster of the relationship.
- Delete the relationship and create the relationship so that the master cluster of the relationship is identical

to the master cluster of the group and the auxiliary cluster of the relationship is identical to the auxiliary cluster of the group.

- Specify a group and a relationship that have identical master clusters and identical auxiliary clusters.

Resubmit the task.

CMMVC6478E The Enable remote authentication service task cannot be initiated because the server settings are not configured.

Explanation: You cannot enable the remote authentication service until the server has been configured with all of the required settings. You must specify the user name, password, and remote authentication server URL, and if required, the SSL certificate.

User response: Ensure that the server settings are configured correctly, and resubmit the task.

CMMVC6479E The task cannot be initiated because the user group table is full.

Explanation: The maximum supported number of user groups is already configured in the user group table.

User response: Delete a user group that is not required from the table, and resubmit the task.

CMMVC6480E The task cannot be initiated because the user group that you have specified is not defined.

Explanation: You must specify a user group that exists in the user group table.

User response: Either create the user group that you had specified or specify an existing user group, and resubmit the task.

CMMVC6481E The Modify user group task cannot be initiated because you have specified a default user group.

Explanation: Examples of default user groups are SecurityAdmin, Administrator, CopyOperator, Service, and Monitor.

User response: Specify a user group that is not a default user group when you submit this task.

CMMVC6482E The Delete user group task cannot be initiated because you have specified a default user group.

Explanation: Examples of default user groups are SecurityAdmin, Administrator, CopyOperator, Service, and Monitor.

User response: Specify a user group that is not a default user group when you submit this task.

CMMVC6483E The task cannot be initiated because the user group name that you have specified already exists.

Explanation: Each user group must have a unique name.

User response: If you want to define a new user group with the name that you had specified, you must first delete the existing user group that has that same name. Specify a user group name that does not exist when you submit this task.

CMMVC6484E The task cannot be initiated because the role that you have specified is not supported.

Explanation: Examples of valid roles are SecurityAdmin, Administrator, CopyOperator, Service, and Monitor.

User response: Specify a supported role, and resubmit the task.

CMMVC6485E The Delete user group task has failed because there is at least one user that is defined as a member of the group, and you did not specify the `-force` parameter.

Explanation: You cannot delete a user group that is not empty unless you specify the `-force` parameter. If you use the `-force` parameter when you delete a user group, all of the users that were in the deleted user group are added to the Monitor user group.

User response: Ensure that you specify the correct user group. For each member of the specified user group that you want to belong to a user group other than Monitor, move that member to the desired group. If the user group has at least one member, specify the `-force` parameter when you submit the task.

CMMVC6486E The task cannot be initiated because the user table is full.

Explanation: The maximum supported number of users is already configured in the user table.

User response: Delete a user that is not required from the table, and resubmit the task.

CMMVC6487E The task cannot be initiated because the user name that you have specified already exists.

Explanation: Each user must have a unique name.

User response: If you want to define a new user with the name that you had specified, you must first delete the existing user that has that same name. Specify a user name that does not exist when you submit this task.

CMMVC6488E The task cannot be initiated because you have specified a user group ID that is not correct.

Explanation: You must specify a valid user group ID when you submit this task.

User response: Specify a valid user group ID, and resubmit the task.

CMMVC6489E The task cannot be initiated because you have specified more than one password.

Explanation: This task allows you to specify only one password.

User response: Specify only one password, and resubmit the task.

CMMVC6490E The task cannot be initiated because you have specified both a user group and the use of the remote authentication service.

Explanation: You cannot specify a user group when you specify the use of the remote authentication service.

User response: Either specify a user group or specify the use of the remote authentication service, but not both, and resubmit the task.

CMMVC6491E The task cannot be initiated because you have specified the use of the remote authentication service but you have not specified an SSH key and password.

Explanation: An SSH key and password are both required to use the remote authentication service.

User response: Specify a valid SSH key and a valid password when you submit this task.

CMMVC6492E The task cannot be initiated because you have specified a local user but you have not specified a user group.

Explanation: You must specify a user group when you specify a local user for this task.

User response: Specify a valid user group if you specify a local user when you submit this task.

CMMVC6493E The task cannot be initiated because the user that you have specified is not defined.

Explanation: You must specify a user that exists in the user table.

User response: Either create the user that you had

specified or specify an existing user, and resubmit the task.

CMMVC6494E The task cannot be initiated because you cannot remove a default user.

Explanation: Examples of default users are SecurityAdmin, Administrator, CopyOperator, Service, and Monitor.

User response: Specify a user that is not a default user when you submit this task.

CMMVC6495E The task cannot be initiated because the user superuser must be a local user.

Explanation: You cannot define the user superuser to use the remote authentication service.

User response: Ensure that you have specified the correct user, and resubmit the task.

CMMVC6496E The task cannot be initiated because you cannot remove the superuser password.

Explanation: The user superuser must always have a password defined.

User response: Ensure that you have specified the correct user when you submit the task.

CMMVC6497E The task cannot be initiated because the user that you have specified does not have a password defined.

Explanation: You cannot remove a password that does not exist.

User response: Ensure that you have specified the correct user when you submit the task.

CMMVC6498E The task cannot be initiated because the user that you have specified does not have an SSH key defined.

Explanation: You cannot remove an SSH key that does not exist.

User response: Ensure that you have specified the correct user when you submit the task.

CMMVC6499E The task has failed because the SSH key that you have specified is already defined for another user.

Explanation: A single SSH key cannot be defined for more than one user.

User response: Either specify a unique SSH key for the user that you had specified or delete the user that has the SSH key that you had specified, and resubmit the task.

CMMVC6500E The action failed because the source and destination virtual disks (VDisks) are the same.

Explanation: The action failed because the source and destination volumes are the same.

User response:

CMMVC6501E The action failed because the node hardware is incompatible with the current I/O group member.

Explanation: The action failed because the node hardware is incompatible with the current I/O group member.

User response:

CMMVC6502E The FlashCopy mapping was not prepared because preparing consistency group 0 is not a valid operation.

Explanation: The FlashCopy mapping was not prepared because preparing consistency group 0 is not a valid operation.

User response:

CMMVC6503E The FlashCopy mapping or consistency group was not stopped because stopping consistency group 0 is not a valid operation.

Explanation: The FlashCopy mapping or consistency group was not stopped because stopping consistency group 0 is not a valid operation.

User response:

CMMVC6504E The task cannot be initiated because the SSH key file that you have specified does not contain a valid SSH key.

Explanation: You must specify an SSH key file that contains a valid SSH key.

User response: Specify an SSH key file that contains a valid SSH key, and resubmit the task.

CMMVC6505E The task has failed because an error has occurred while communicating with the authentication service.

Explanation: The cluster is configured to use an authentication service to control which users are authorized to access the cluster. An error has occurred while the cluster was attempting to contact the authentication service. The error is probably the result of an incorrect configuration, either of the cluster or of the authentication service. This error occurs if the SSL certificate, user name or password is incorrect.

CMMVC6506E • CMMVC6511E

User response: Ensure that the authentication service is functioning properly. Ensure that the cluster authentication service configuration is correct. Resubmit the task.

CMMVC6506E The task has failed because a timeout has occurred while communicating with the authentication service.

Explanation: The cluster is configured to use an authentication service to control which users are authorized to access the cluster. A timeout has occurred while the cluster was attempting to contact the authentication service. This timeout is probably the result of a TCP/IP network problem or of an incorrect configuration. Configuring the incorrect IP address or protocol in the authentication service URL causes this error. The protocol can be either http or https.

User response: Ensure that the cluster authentication service configuration is correct. Ensure that the Ethernet network between the cluster and the authentication service is functioning properly. Ensure that the authentication service is functioning properly. Resubmit the task.

CMMVC6507E The task has failed because the authentication service reports an incorrect user name or password.

Explanation: The cluster is configured to use an authentication service to control which users are authorized to access the cluster.

If the password for the user name has recently been changed on the authentication service, you might be required to force the cluster to refresh its authentication cache. You can force the refresh using the cluster console View Cluster Properties, Remote Authentication panel or by submitting the Command-Line Interface command `svctask chauthservice -refresh`.

User response: Ensure that the user name and password that you use is correct.

If the password for the user name has recently been changed on the authentication service, force the cluster to refresh its authentication cache.

If the user name that you are using also has a password configured on the cluster, ensure that the password that is configured on the cluster is identical to the password that is configured for that user name on the authentication service.

Resubmit the task.

CMMVC6508E The task has failed because the authentication service reports that the authentication token has expired.

Explanation: The cluster is configured to use an authentication service to control which users are authorized to access the cluster. The authentication

token, which is saved as a browser cookie, has expired. You can modify the token expiration property that is set by the authentication service to reduce the frequency of this error in the future.

User response: Either acquire a new authentication token or log in with a user name and password, and resubmit the task.

CMMVC6509E The task has failed because the user name is not configured on the cluster.

Explanation: If the user name is defined on an authentication service and you want to use that service for cluster authentication, you must configure the cluster to use that authentication service.

User response: Ensure that you are using the correct user name.

If the user name is not configured on the cluster and you want to use the cluster to authenticate, create a new user with the user name that you want to use on the cluster.

If the user name is defined on an authentication service and you want to use that service for cluster authentication, configure the cluster to use that authentication service.

Resubmit the task.

CMMVC6510E The task has failed because the password is not correct for the user name that you are using.

Explanation: The password that you are using does not match the password that is configured on the cluster for the user name that you are using.

User response: Enter the correct password, and resubmit the task.

CMMVC6511E The task has failed because the cluster is not configured correctly to use the authentication service.

Explanation: The user name that you are using is configured to be authenticated using an authentication service, but either the cluster is not configured to use an authentication service or the function is not enabled.

User response: If you want to use an authentication service, configure the cluster to use the service.

If you do not want to use an authentication service, modify the configuration of the user name on the cluster to remove the designation for the use of the authentication service.

Resubmit the task.

CMMVC6512E The task has failed because you cannot both create a new quorum disk and set that new disk to active using the same command.

Explanation: The create new quorum disk task and set disk to active task must be done using two separate tasks.

User response: Submit a create new quorum disk task. When that task has completed, submit a task to activate the new disk.

CMMVC6513E The task has failed because you cannot activate a quorum disk until all of the quorum disks have been initialized.

Explanation: The initialization process for at least one disk has not yet completed. You cannot select a disk as the active disk until the initialize process for all of the quorum disks has completed.

User response: Wait until the initialize quorum disk process has completed for all of the quorum disks, and resubmit the task.

CMMVC6514E The task has failed because the disk that you have selected to activate is not online.

Explanation: A disk must be online to be eligible for activation.

User response: Either bring the disk that you have selected online or select a different disk that is already online, and resubmit the task.

CMMVC6515E The task has failed because at least one quorum disk is in the Excluded state.

Explanation: You cannot activate a quorum disk when one or more of the quorum disks are in the Excluded state.

User response: Either create additional quorum disks or change the configuration so that none of the quorum disks is in the Excluded state. Ensure that none of the quorum disks are in the Excluded state, and resubmit the task.

CMMVC6516E The command has failed because an IPv4 cluster address cannot be removed while remote IPv4 services are configured.

Explanation: The configured cluster IP address protocols determine whether IPv4 or IPv6 or both are enabled on the cluster. If a cluster does not have an IPv4 cluster address the IPv4 protocol stack will not be enabled, and therefore remote services such as email

servers or SNMP servers cannot be accessed through an IPv4 address.

User response: If you can only access the service through an IPv4 address and you need to continue to use the service, you will also have to continue to specify an IPv4 cluster address even if you do not intend to manage your cluster through this address.

Otherwise, re-configure the cluster so that all remote services use only IPv6 addresses, and resubmit the task to remove the IPv4 cluster address.

CMMVC6517E The command has failed because an IPv6 cluster address cannot be removed while remote IPv6 services are configured.

Explanation: The configured cluster IP address protocols determine whether IPv4 or IPv6 or both are enabled on the cluster. If a cluster does not have an IPv6 cluster address the IPv6 protocol stack will not be enabled, and therefore remote services such as email servers or SNMP servers cannot be accessed through an IPv6 address.

User response: If you can only access the service through an IPv6 address and you need to continue to use the service, you will also have to continue to specify an IPv6 cluster address even if you do not intend to manage your cluster through this address.

Otherwise, re-configure the cluster so that all remote services use only IPv4 addresses, and resubmit the task to remove the IPv6 cluster address.

CMMVC6518E The task has failed because no roles are defined for the current user on the cluster.

Explanation: The cluster has been configured to use an authentication service to control which users are authorized to access the cluster. The user's credentials were accepted by the authentication service, but none of the groups defined for the user on the authentication service match any of the user groups that are defined on the cluster.

User response: Perform the following steps in sequence:

1. Determine which user groups are defined for the user on the authentication service.
2. Ensure that at least one user group that is defined for the user on the authentication service is also defined on the cluster.
3. Ensure that at least one user group that is defined for the user on both the authentication service and the cluster has its 'remote' parameter set to 'enabled'.
4. Resubmit the task.

CMMVC6519E The task has failed because you cannot change the user group of the 'superuser' account to anything other than 'SecurityAdmin'.

Explanation: The user group that is assigned to the user name 'superuser' must always be 'SecurityAdmin'. This assignment cannot be changed.

User response: Ensure that you specify a user account other than 'superuser' if you submit a task to change the user group of a user account from 'SecurityAdmin' to a different user group.

CMMVC6520E You cannot use this task to modify the properties of the current user because those properties are only defined by an authentication service.

Explanation: The current user is not defined on the cluster. The current user is defined on an authentication service, and the cluster is configured to use that authentication service. You must use the authentication service to change the current user's password.

If you want to enable command-line interface (CLI) access to the cluster by using an SSH key, you must define the current user on the cluster and associate the SSH key with that user. If you also want to continue using the authentication service for the current user, you must enable the 'remote' setting for the new current user account that you create on the cluster.

User response: If you want to change your password, use the authentication service for that task.

If you want to enable command-line interface (CLI) access to the cluster by using an SSH key, define your user account on the cluster and associate the ssh key with that definition. If you also want to continue using the authentication service to authorize your user account, enable the 'remote' setting for your newly created user account on the cluster.

CMMVC6521E The task cannot be initiated because it would have resulted in a user account definition for a local user that specifies neither a password nor an SSH key.

Explanation: The definition of a local user must always specify either a password or an SSH key.

User response: When you submit this task, ensure that you have specified the correct user account and parameters, and that all local user definitions would still specify either a password or an SSH key after the task completes.

CMMVC6522E Authorization has failed.

Explanation: An SSH login attempt has failed. This message will be followed by a second message that will contain detailed information about the cause of the error.

User response: Follow the instructions in the second error message to solve the problem.

CMMVC6523E The URL that you have entered is not valid.

Explanation: The URL must start with either http:// or https:// and must use only the following characters: A-Z, a-z, 0-9, - _ : [] . ~ / %.

User response: Ensure that the URL that you enter starts with one of the supported strings and contains only supported characters, and resubmit the task.

CMMVC6524E The name that you have entered is not valid. The name cannot begin or end with a space character, and the name cannot contain any of the following characters: * , \ ' %

Explanation: A space cannot be the first or last character in the name that you enter. Also, the following characters are not supported anywhere in the name: * , \ ' %

User response: Ensure that the name that you enter does not begin or end with the space character and that it does not contain any of the unsupported characters listed above, and resubmit the task.

CMMVC6525E The password that you have entered is not valid. The password cannot begin or end with a space character.

Explanation: A space cannot be the first or last character in the password that you enter.

User response: Ensure that the password that you enter does not begin or end with the space character, and resubmit the task.

CMMVC6526E The Create VDisk task cannot be initiated because the number of copies that you have requested is not equal to the number of unique MDisk groups that you have specified.

Explanation: When you submit this task, you must specify a unique storage pool for each volume copy that you request.

User response: Specify the same number of unique storage pools as the number of volume copies that you request, and resubmit the task.

CMMVC6527E The name that you have entered is not valid. The name can contain letters, numbers, spaces, periods, dashes, and underscores. The name must begin with a letter or an underscore. The name must not begin or end with a space.

Explanation: A number or space cannot be the first character and a space cannot be the last character in the name that you enter. Also, the following characters are not supported anywhere in the name: * : , " ' % #

User response: Ensure that the name you enter does not begin with a number, does not begin or end with a space character, and does not contain any of the unsupported characters listed above, and resubmit the task.

CMMVC6528E The command cannot be initiated because the MDisk mode is not set to Array.

Explanation: Any MDisk that you specify for this command must be a local MDisk that is an array of LDisks. The mode of the MDisk that you have specified is not Array.

User response: Either select a different MDisk that is a local MDisk and is an array of LDisks, or configure the system so that the MDisk that you have specified is a local MDisk and is an array of LDisks, and resubmit the command.

CMMVC6529E The command cannot be initiated because the maximum supported number of MDisks already exists.

Explanation: This command requires that an MDisk is available for array creation. There are no available MDisks for array creation because the maximum number of MDisks is already configured on the cluster.

User response: Ensure that a local MDisk is available, and resubmit the command. To make a local MDisk available for this task, either delete an array on an existing local MDisk or remove a SAN attached MDisk and configure a local MDisk.

CMMVC6530E The command cannot be initiated because the maximum supported number of arrays already exists.

Explanation: The cluster already has the maximum number of arrays that it can support. The command attempted to add a new array.

User response: Remove an array that is no longer needed, and resubmit the command.

CMMVC6532E The command cannot be initiated because there is insufficient free memory that is available to the I/O group.

Explanation: This command requires that there is sufficient free memory available for the specified I/O group to allocate the memory that is required for the new array.

User response: Ensure that there is sufficient memory available to the I/O group, and resubmit the command. You can increase the amount of memory that is allocated to the I/O group. You can also reduce the amount of memory that is used by reducing the number volume mirrors or Copy Services relationships in the I/O group.

CMMVC6533E The command cannot be initiated because the specified array member does not exist in the selected array.

Explanation: This command requires that the array member that you specify is an LDisk. It is possible that the array member that you specified was an LDisk that was recently deconfigured due to an error. You can use the **lsarraymember** command to display the available members of an array.

User response: Select an array member that has an associated LDisk, and resubmit the command.

CMMVC6534E The command cannot be initiated because the drive that you have specified does not exist.

Explanation: You have specified a drive ID that is not defined.

User response: Use the **svcinfo lsdrive** command to display existing drive IDs. Specify only existing drive IDs, and resubmit the command.

CMMVC6535E The command cannot be initiated because you have specified an insufficient number of drives to configure an array using the RAID geometry that you have specified.

Explanation: Each RAID geometry requires a minimum number of available drives in order to configure an array using that geometry. For example, a RAID 6 geometry requires that you specify at least four available drives. The number of drives that you have specified is less than the minimum number of drives that are required for the RAID geometry that you have specified.

User response: Ensure that you specify a sufficient number of drives to accommodate the RAID geometry that you specify, and resubmit the command. You might want to specify a different number of drives or a different RAID geometry.

CMMVC6536E The command cannot be initiated because you have specified more drives than the specified RAID geometry permits.

Explanation: The number of drives that you specify must be within the supported range of number of drives that is supported for the RAID geometry that you specify. For example, a RAID 1 geometry requires that you specify exactly two available drives.

User response: Specify a number of available drives that is supported for the RAID geometry that you specify, and resubmit the command.

CMMVC6537E The command cannot be initiated because the drive that you have specified has a Use property that is not supported for the task.

Explanation: You can submit the `svcinfo lsdrive` command to display the Use property of a drive and to determine which drives are available.

User response: Consult the command documentation to determine what drive Use property values are supported for this command. Ensure that you select a drive that has a value for the Use property that is supported when you submit this command.

CMMVC6538E The command cannot be initiated because at least one of the drives that you have specified has a Use property that is not Candidate.

Explanation: Every drive that you specify for this command must have a Use property of Candidate. You can submit the `svcinfo lsdrive` command to display the Use property of existing drives.

User response: Ensure that all of the drives that you specify have a Use property of Candidate, and resubmit the command.

CMMVC6539E The command cannot be initiated because the array does not have sufficient redundancy.

Explanation: The array must have sufficient redundancy when you submit this command. The task that you have requested would have taken the array offline.

User response: Fix all errors that are related to the array that you have specified and restore redundancy to the array before you resubmit the command.

CMMVC6540E The task cannot be initiated because the space-efficient grain size is too small to accommodate the virtual capacity that you have requested for the VDisk.

Explanation: The virtual capacity that you have

requested would require a larger number of grains than the supported maximum for the specified grain size.

User response: Either increase the grain size, decrease the requested virtual capacity of the volume, or both, and resubmit the task.

CMMVC6541E The task cannot be initiated because the virtual capacity that you have requested for the space-efficient VDisk is larger than the maximum capacity that is supported for the extent size.

Explanation: The extent size of the storage pool that you have selected would require a larger number of extents than the supported maximum to accommodate the virtual capacity that you have requested for the thin-provisioned volume.

User response: Either select a different storage pool that has an extent size that is large enough to accommodate the requested virtual capacity or specify a virtual capacity that is supported for the extent size of the storage pool that you had selected, and resubmit the task.

CMMVC6542E The remote authentication task has failed.

Explanation: An error has occurred while attempting to authenticate a user account using a remote authentication service. You can run the `svc_snap` task to gather cluster information that can be used in problem determination.

User response: Contact IBM technical support for assistance.

CMMVC6543E The task cannot be initiated because you can only specify a direct-attached managed drive when you submit the task.

Explanation: The drive that you have specified either is not managed or is not a local drive.

User response: Specify a direct-attached MDisk when you submit this task.

CMMVC6544E The task cannot be initiated at this time because the direct-attached managed drive that you have specified is too busy. Resubmit the task when the drive is less busy.

Explanation: The task takes approximately thirty seconds to complete. When the direct-attached managed drive is busy, the time that is required to complete the task increases. When the drive is too busy, the task cannot complete in a reasonable amount of time.

User response: Resubmit the task when the direct-attached managed drive is less busy.

CMMVC6545E The Apply Drive Software task has failed to access the software download image.

Explanation: Either the image file cannot be read, the validation signature is incorrect, the drive type or firmware type is not correct, or the image file has been corrupted.

User response: Reinstall the firmware download image, and resubmit the task. If the problem persists, contact IBM technical support for assistance.

CMMVC6546E A device error was detected during the Apply Drive Software task.

Explanation: The task might have succeeded.

User response: View the event in the event log. Determine the firmware level from the VPD of the node. If the VPD does not show that the downloaded firmware version is installed, resubmit the task.

CMMVC6547W The Download FPGA firmware task has been initiated. The MDisk remains Offline while the task is in progress. Do not remove power from the drive or node while the task is in progress.

Explanation: The task might take approximately fifteen minutes to complete. When the task completes, the drive status changes to Online automatically.

User response: Ensure that electrical power is continuously being supplied to the node and the drive, at least until the task completes and the drive status changes to Online.

CMMVC6548E The FPGA firmware cannot be applied because the drive has a use other than candidate.

Explanation: Updating a drive FPGA level is not guaranteed to maintain data integrity, therefore the drive must not be part of an array. To ensure this, the drive must have a use of "candidate" before the package can be applied.

User response: If the drive is currently in the "failed" state, run through all maintenance actions required for the drive before continuing. If the drive is a spare or unused, the drive use can be changed through the GUI or through the **chdrive** command. If a drive is currently part of an array, a hot spare must be configured and the drive use changed to failed, before changing the use to candidate.

CMMVC6549E The Authentication task has failed because the authentication service URL that you have specified is not a valid URL.

Explanation: This error might be caused by the authentication service not operating correctly or by an incorrect URL being defined for the authentication service. You can use the **svctask chauthservice** command to change the URL that is defined in the cluster for the authentication service.

User response: Ensure that the authentication service is operating correctly. Ensure that the authentication service URL that is defined in the cluster is correct, and resubmit the task.

CMMVC6550E The Authentication task has failed because the network address that is specified in the authentication service URL cannot be resolved.

Explanation: The authentication service URL that is defined in the cluster has a network address that cannot be resolved. You can use the **svctask chauthservice** command to change the URL that is defined in the cluster for the authentication service.

User response: Ensure that the authentication service is operating correctly. Ensure that the authentication service URL that is defined in the cluster is correct. Ensure that the network connection between the cluster and the authentication service is operating correctly, and resubmit the task.

CMMVC6551E The Authentication task has failed because the combination of user name and password that is defined in the cluster for authorization by the authentication service is not defined on the authentication service.

Explanation: The authentication service has refused an authentication request from the cluster. You can use the **svctask chauthservice** command to change the user name or the password that is defined in the cluster for the authentication service.

User response: Ensure that the user name and password combination that is defined in the cluster for the authentication service is also defined on the authentication service, and resubmit the task.

CMMVC6552E The Authentication task has failed because an SSL connection could not be established with the authentication service.

Explanation: This error might be caused by an incorrect SSL configuration on the authentication service server or by a rejection by the authentication service server of the SSL certificate that is configured

on the cluster. You can use the `svctask chauthservice` command to set the SSL certificate that is defined in the cluster for the authentication service server.

User response: Ensure that the SSL configuration on the authentication service server is correct and that the SSL certificate that is defined in the cluster for the authentication service server is correct, and resubmit the task.

CMMVC6553E The task cannot be initiated because at least one quorum disk is not in the correct state.

Explanation: All of the quorum disks must have a state of Online when you set an MDisk to be the active quorum disk.

User response: Ensure that all of the quorum disks have a state of Online, and resubmit the task.

CMMVC6554E The Authentication task has failed because the user name that was received from the authentication service is not a valid cluster user name.

Explanation: The cluster user name cannot exceed 256 characters in length, and cannot contain any of the following characters:

- colon :
- percent sign %
- comma ,
- double quote ""
- single quote '

User response: Change the definition of the user name in the remote authentication service so that it conforms to the cluster user name requirements, and resubmit the task.

CMMVC6555E The Authentication task has failed because the authentication service either sent an incorrect response, or it sent a response that indicates that the authentication request has failed for a reason other than incorrect authentication credentials.

Explanation: Either the format of the response from the authentication service is not valid or the response indicates a failure to authenticate that is not related to the credentials that were being authenticated.

User response: Ensure that the authentication service is functioning correctly, and resubmit the task. If the problem persists, contact the authentication service technical support for assistance.

CMMVC6556E The task cannot be initiated because an error has occurred while attempting to read a file.

Explanation: The task specified the name of a file on the file system of the cluster configuration node. The specified file cannot be opened. This error might be caused by a typographical error in the file name that you specified or by a failover of the configuration node to a different node than the node into which you are currently logged in.

User response: Ensure that the file has been copied to the current configuration node and that you are logged in to that node, specify the correct file name, and resubmit the task.

CMMVC6557E The task cannot be initiated because the file that you have specified is too large.

Explanation: The task specified the name of a file on the file system of the cluster configuration node. The specified file cannot be used because it exceeds the maximum size supported for the task. If the file has been corrupted, you can copy the correct version of the file onto the configuration node to restore the correct file size. The maximum file size is described in the task help.

User response: Specify the correct file name and ensure that the size of the file does not exceed the supported maximum file size for this task, and resubmit the task.

CMMVC6558E The command cannot be initiated because it might cause VDisks to go Offline. Refer to the cluster Command-Line Interface (CLI) command help for this command.

Explanation: You are being warned that this command might result in taking volumes Offline. After you completely understand the possible consequences by reading the command help, you can override the safety precautions and avoid this message by using the `-force` flag.

User response:

1. Submit the `lsnode dependantvdisks` command to determine which volumes will go Offline if you resubmit this command using the `-force` flag. If you received this message when you submitted the `applysoftware` command, you must submit the `lsnode dependantvdisks` command for every node in the cluster; for all other commands you must submit the `lsnode dependantvdisks` command for the node that you specified as a parameter in the command that generated this message.
2. This step is required because it is critically important that you understand the implications of using the `-force` flag for the specific command that

you have submitted: Refer to the CLI command help to determine what safety precautions are bypassed when you use the `-force` flag. The ignored precautions differ, depending on the command.

3. If you want to bypass the safety precautions when you resubmit the command, you must use the `-force` flag.

CMMVC6559E The Add or Change email user command has failed because you have specified a user type of 'support' and you have specified either the `-warning` or `-info` parameter value as 'on'.

Explanation: The user type 'support' is intended to be used to indicate that the user is from a hardware maintenance support service external to your organization. Therefore, only events with the more serious notification type of 'error' can be sent to a 'support' user type.

User response: Ensure that you have specified the correct user type. If you want this user to receive warning or information notifications, do not specify the `-usertype support` parameter and value. If you specify the user type as 'support', you must specify the `-warning` and `-info` parameters as 'off'.

CMMVC6560E The command has failed because the specified IP address is already in use by the cluster.

Explanation: You cannot specify an IP address that is already configured to be used by the cluster.

User response: Ensure that the IP address that you specify is not already configured for use by the cluster, and resubmit the task.

CMMVC6561E The set quorum active task has failed because either another set quorum active task is in progress or the selected disk cannot be set as the active quorum disk.

Explanation: This is a multi-step task and can take from a few seconds to several minutes to complete. Only one set quorum active task can be in progress at any specified time. This error has one of two causes. Either another set quorum task is already in progress, or the internal cluster logic did not accept your request to make the selected disk the active quorum disk.

User response: Check the state of the MDisks and complete any outstanding directed maintenance procedures (DMPs). If another set quorum active task might be in progress, wait for several minutes for that task to complete, and resubmit this task. If you have received this error when there is no other set quorum active task in progress, specify a different disk to replace the current active quorum disk and specify the same quorum index number, and resubmit this task.

CMMVC6562E The requested size exceeds the maximum supported value.

Explanation: You have submitted a command that has a size parameter and an associated unit option that has a default value of Megabytes (MB, 2e20 bytes) when the `-unit` option is not specified. The value that you have specified for the size parameter in combination with the specified or default unit value is greater than the maximum supported size of (2e64 - 1) bytes.

User response: Ensure that the size that you specify is correct for the value of the unit option that is defaulted or specified, and that the size is not greater than the maximum supported size, and resubmit the task.

CMMVC6563E The command has failed because a user that you have specified is not configured to receive email notifications.

Explanation: All of the users that you specify as a target recipient in the `testemail` command must already have at least one of the following email notification flags set to 'on': `-error`, `-warning`, or `-info`.

User response: Ensure that all of the users that you specify have at least one email notification flag set to 'on', and resubmit the command.

CMMVC6564E You cannot make this user a remote user because the password type is not valid for a remote user.

Explanation: The remote authentication server has requirements that do not accept legacy type passwords. This user has a legacy type password.

User response: Either specify a new password and resubmit the command, or first modify the password and then resubmit the command to designate remote authentication for this user.

CMMVC6565E The command has failed because the specified node is offline.

Explanation: This command requires that the status of the node that you specify is Online.

User response: Ensure that the node that you specify has a status of Online when you submit this command.

CMMVC6566E The command cannot be submitted because specifying the `-failover` parameter requires that you also specify either the `-name`, `-iscsialias` or `-noiscsialias` parameter.

Explanation: You have not specified the required failover data that is required when you specify the `-failover` parameter.

User response: Ensure that you want to specify the

CMMVC6567E • CMMVC6574E

-failover parameter. When you specify the -failover parameter with this command, ensure that you also specify either the -name, -iscsialias or -noiscsialias parameter.

CMMVC6567E The Apply Drive Software task has failed because no download images were found in the package file.

Explanation: The drive software upgrade package file was unpacked but no download software images were found in the package.

User response: Acquire a valid solid-state drive software upgrade package file, and resubmit the task using the new package file.

CMMVC6568E The Apply Drive Software task has failed because no download images were found in the package file for this drive type.

Explanation: The package file documentation lists the drive types for which there are images.

User response: Acquire a valid solid-state drive software upgrade package file that contains an image for this drive type, and resubmit the task using the new package file.

CMMVC6569E The Apply Drive Software task has failed because no download images were found in the package file of this software type.

Explanation: The package file documentation lists the drive types and software types for which there are images. The value of the -type parameter that you enter for software type is case-sensitive.

User response: Ensure that the value that you enter for the -type parameter exactly matches the software type that is contained in the solid-state drive software upgrade package file, and resubmit the command.

CMMVC6570E The command was not initiated because the cache mode of the virtual disk (VDisk) is already in the state that you had requested.

Explanation: You have issued a change volume cache mode command but requested the current mode, so there would not have been a change. Therefore, the command was ignored.

User response: List the volume properties to determine the current cache mode. If you want to change the cache mode, ensure that you specify a cache mode that is different from the current cache mode, and resubmit the command.

CMMVC6571E The command has failed because the I/O group that manages the virtual disk (VDisk) that you specified was offline when you submitted the command. You can use the -force flag to force the operation, which might result in the loss of cache data.

Explanation: If you submit this command without the -force flag, the I/O group that manages the volume that you specify must have a state of Online.

Note: Use of the -force flag when you change the cache mode might result in loss of the cache data for the volume, depending on the current cache mode and requested cache mode. One example of a risk of potential loss of cache data would be changing the cache mode from readwrite to none.

User response: Either follow service procedures to bring the I/O group online or specify the -force flag to force the change of the cache mode of the volume, and resubmit the task.

CMMVC6572E The command has failed because the I/O group that manages the virtual disk (VDisk) that you specified is not stable.

Explanation: The unstable I/O group condition is typically transient, and usually occurs during I/O group failover or fail back processing.

User response: Wait a few minutes, and resubmit the command.

CMMVC6573E The command has failed because the VDisk that you specified is a source or target of a FlashCopy mapping that is in the prepared state.

Explanation: If the volume is the source or target of a FlashCopy mapping, the FlashCopy mapping must be in the idle_copied state or the stopped state when you change the cache mode of the volume.

User response: Either remove or stop the FlashCopy mapping and wait for the FlashCopy mapping state to become idle_copied or stopped, and resubmit the command.

CMMVC6574E The command has failed because the VDisk that you specified is a source or target of a FlashCopy mapping that is in the suspended state.

Explanation: If the volume is the source or target of a FlashCopy mapping, the FlashCopy mapping must be in the idle_copied state or the stopped state when you change the cache mode of the volume.

User response: Either remove or stop the FlashCopy mapping and wait for the FlashCopy mapping state to

become `idle_copied` or `stopped`, and resubmit the command.

CMMVC6575E The command has failed because the VDisk that you specified is a source or target of a FlashCopy mapping that is in the preparing state.

Explanation: If the volume is the source or target of a FlashCopy mapping, the FlashCopy mapping must be in the `idle_copied` state or the `stopped` state when you change the cache mode of the volume.

User response: Either remove or stop the FlashCopy mapping and wait for the FlashCopy mapping state to become `idle_copied` or `stopped`, and resubmit the command.

CMMVC6576E The command has failed because the VDisk that you specified is a source or target of a FlashCopy mapping that is in the stopping state.

Explanation: If the volume is the source or target of a FlashCopy mapping, the FlashCopy mapping must be in the `idle_copied` state or the `stopped` state when you change the cache mode of the volume.

User response: Either remove or stop the FlashCopy mapping and wait for the FlashCopy mapping state to become `idle_copied` or `stopped`, and resubmit the command.

CMMVC6577E The command has failed because the VDisk that you specified is a source or target of a FlashCopy mapping that is in the copying state.

Explanation: If the volume is the source or target of a FlashCopy mapping, the FlashCopy mapping must be in the `idle_copied` state or the `stopped` state when you change the cache mode of the volume.

User response: Either remove or stop the FlashCopy mapping and wait for the FlashCopy mapping state to become `idle_copied` or `stopped`, and resubmit the command.

CMMVC6578E The command has failed because the iSCSI name is already assigned or is not valid.

Explanation: The cluster does not support duplicate iSCSI names. A valid iSCSI name cannot contain a comma or leading or trailing spaces.

User response: Ensure that you specify a unique and valid iSCSI name, and resubmit the command.

CMMVC6579E The command cannot be initiated because the cluster Ethernet port 1 must always be fully configured in either the IPv4 or IPv6 format.

Explanation: This error can be caused by an attempt to delete the only address that is configured on the primary Ethernet port on the cluster.

User response: When you delete an IP address on the primary Ethernet port, ensure that the other supported IP format is already configured on that port.

CMMVC6580E The command cannot be initiated because the iSCSI alias that you specified contained either leading or trailing space characters.

Explanation: The space character cannot be the starting or ending character of an iSCSI alias name.

User response: Ensure that the iSCSI alias that you specify does not begin or end with a space character, and resubmit the command.

CMMVC6581E The command has failed because the maximum number of allowed iSCSI qualified names (IQNs) has been reached, or the IQN is already assigned or is not valid.

Explanation: IQNs cannot exceed the maximum number allowed, must not be duplicated, must not contain a comma, and must not contain leading or trailing spaces.

User response: If the number of IQNs is within the allowed maximum, ensure that you specify a unique and valid IQN, and resubmit the command.

CMMVC6582E The task has failed because the iSCSI host that you specified is not mapped to an I/O group.

Explanation: You cannot add a port to an iSCSI host until you have mapped the iSCSI host to at least one I/O group.

User response: Map the iSCSI host to at least one I/O group, and resubmit the command.

CMMVC6583E The command has failed because the name that you specified contains a character that is not supported for a node or cluster name.

Explanation: A node or cluster name cannot contain any of the following characters or ASCII hexadecimal values:

- 0000-001F ASCII control characters
- 0020-002C The space character `!""# $ % the ampersand character ' () * + ,`

CMMVC6584E • CMMVC6591E

- 002F /
- 003B-0040 ; the 'less than' character = > ? @
- 005B-0060 [\] ^ _ `
- 007B-007F { | } ~ and the DEL character

User response: Specify a valid name, and resubmit the command.

CMMVC6584E The command cannot be initiated because it would deconfigure the remote authentication service while the service is enabled.

Explanation: You cannot deconfigure the remote authentication service while it is enabled.

User response: Ensure that the remote authentication service is not being used, disable the service, and resubmit the task.

CMMVC6585E The command cannot be initiated because the array that you have specified has a geometry of RAID 0, which is not a redundant geometry.

Explanation: The array that you specify for this command must have a redundant geometry, and RAID 0 is not a redundant geometry.

User response: Ensure that you specify an array that has a redundant geometry when you submit the command.

CMMVC6586E The command cannot be initiated because the action would cause array data loss due to the unsynchronized state of the array.

Explanation: To avoid data loss, this command is not permitted to process an array that is not synchronized.

User response: Use the `lsarraysyncprogress` command to ensure that the synchronization process completes for this array, and resubmit the task.

CMMVC6587E The command did not complete because I/O for the array was not quiesced within the allotted time period.

Explanation: All outstanding I/O for the array must complete before the configuration can be changed. The command has failed because there is still outstanding I/O to be processed for the array, and the maximum amount of time allotted to the command has expired.

User response: Resubmit the task.

CMMVC6588E The command cannot be initiated because a drive that you have specified has a capacity that is smaller than the minimum capacity that is required for the array that you have specified.

Explanation: You can use the `svcinfolarraymembergoals` command to identify the capacity requirement for a member of the array that you specified.

User response: Specify a drive that has sufficient capacity for the array that you specify when you submit the command.

CMMVC6589E The command was not initiated because the drive that you have specified does not sufficiently match the array member goals and you did not specify the `-balanced` parameter.

Explanation: If you do not specify the `-balanced` parameter, the new drive must be an exact match for the array member goals when you exchange a new drive for an existing array member. The new drive that you have specified does not match the goals. If you want to use the drive that you specified to replace an existing member of the array that you specified, you must specify the `-balanced` parameter, which forces the array member goals to be changed to accommodate the new drive.

User response: Either select a different drive that matches the array member goals or specify the `-balanced` parameter to force a change in the array member goals to accommodate the new drive, and resubmit the command.

CMMVC6590E The command cannot be initiated because you did not specify the `-allowdegraded` parameter and the associated array member has insufficient spare protection.

Explanation: This command requires that spare drives are available to assume the function of array member drives that are removed from the array. The requirement can be bypassed by using the `-allowdegraded` parameter.

User response: Either configure sufficient additional spare drives or specify the `-allowdegraded` parameter, and resubmit the command.

CMMVC6591E The command cannot be initiated because the specified sequence number does not match the sequence number of any errors in the error log.

Explanation: The sequence number that is specified in the command must be identical to the sequence number of an event in the event log.

User response: Check the event log to verify the sequence number of the event that you want to specify, and resubmit the command using the correct sequence number.

CMMVC6592E The command cannot be initiated because at least one parameter that was specified is not supported when submitting a command to view the details of an error log entry.

Explanation: Filtering parameters such as '-order severity' or '-status alert' that are valid when listing multiple event log entries are not supported for the command to view the details of a single event log entry.

User response: Check the command syntax, and use supported syntax when you submit the command.

CMMVC6593E The command cannot be initiated because the error log entry has a status that is not supported for the command.

Explanation: Only events with a status of 'alert' or 'message' can be manually marked as fixed or unfixed. Events with a status of 'monitoring' or 'expired' are not required to be marked as fixed or unfixed.

User response: Check the event log to verify the sequence number of the event that you want to specify. Ensure that the event that you specify has a status that is supported for the command when you submit the command.

CMMVC6594E The command cannot be initiated because a drive was specified twice in the drive list.

Explanation: The drive list cannot contain any duplicate entries because the same drive cannot be a member of an array more than once.

User response: Ensure that the drive list that you specify does not contain any duplicate entries when you submit this task.

CMMVC6595E The command cannot be initiated because a drive that you have specified has a technology type that is not supported for the command.

Explanation: The command supports only certain drive technology types. You have specified at least one drive that has a technology type that is not supported for the command.

User response: Consult the command documentation to determine which drive technology types are supported for the command. Submit the **svcinfo lsdrive** command to determine which drives are available. Specify an available drive that has a technology type that is supported for the command

when you submit the command.

CMMVC6596E The command has failed because you have specified an I/O group that does not exist.

Explanation: You must specify an existing I/O group when you submit this command.

User response: Specify an existing I/O group, and resubmit the command.

CMMVC6597E The command has failed because the email settings are not configured.

Explanation: The email system settings must be configured before you can submit a command for error notifications.

User response: Configure the email system settings to enable error notifications, and resubmit the command.

CMMVC6608E The command cannot be initiated because Easy Tier is active on the virtual disk copy.

Explanation: Easy Tier is active on the volume copy, which prevents the success of the command.

User response: Disable Easy Tier on the volume copy or on the storage pool in which the volume copy resides, and resubmit the command.

CMMVC6609E The command cannot be initiated because the size of the Mdisk is smaller than the extent size of the MDisk group.

Explanation: The sizing of the Mdisk in relation to the storage pool is not correct, which prevents the success of the command.

User response: Use a larger Mdisk or make the extent size of the storage pool smaller than the Mdisk, and resubmit the command.

CMMVC6610E The software upgrade cannot start because one or more I/O groups are in maintenance mode.

Explanation: Maintenance mode is used during system servicing, which prevents software upgrades.

User response: Complete system servicing, turn off maintenance mode, and resubmit the command.

CMMVC6611E The command has failed because the specified enclosure is offline.

Explanation: The specified enclosure is offline, which has prevented the success of the command.

User response: Fix any errors associated with the specified enclosure, and resubmit the command.

CMMVC6612E The command has failed because of a hardware error.

Explanation: A hardware error has occurred, which has prevented the success of the command.

User response: Fix any errors in the specified object, and resubmit the command.

CMMVC6613E The command has failed because the specified enclosure type is not supported.

Explanation: You have attempted to use an enclosure of an unsupported type.

User response: Do not attempt to use the specified enclosure type.

CMMVC6614E The command has failed because the specified canister is offline.

Explanation: The specified canister is offline, which has prevented the success of the command.

User response: Fix any errors associated with the specified canister, and resubmit the command.

CMMVC6615E The command cannot be initiated because the selected node is in a different enclosure.

Explanation: You have attempted to add a second node to an I/O group that is not in the same enclosure as the first node. The two nodes of an I/O group must be in the same enclosure.

User response: Select a different node from the list of candidate nodes, and resubmit the command.

CMMVC6616E All available quorum disks are dependent on the MDisks that you have specified.

Explanation: The list of MDisks that you have specified contains all activated quorum disks. If all of the MDisks in the list were to become inaccessible, the system would be unable to backup important data. Operating the system without any online quorum disks is not recommended.

User response: Move one or more quorum disks to MDisks that will remain online.

CMMVC6617E All available quorum disks are dependent on the drives that you have specified.

Explanation: The list of drives that you have specified contains all activated quorum disks. If all of the drives in the list were to become inaccessible, the system would be unable to backup important data. Operating

the system without any online quorum disks is not recommended.

User response: Move one or more quorum disks to drives that will remain online.

CMMVC6618E All available quorum disks are dependent on the enclosure that you have specified.

Explanation: Before removing the enclosure that you have specified, the system must be configured so at least one of the drives that are allocated to hold quorum will remain online when the enclosure goes offline.

User response: Assign one or more drives in the control enclosure as a quorum drive. After you have configured the quorum drives, test for dependencies.

CMMVC6619E All available quorum disks are dependent on the canister that you have specified.

Explanation: Before removing the canister that you have specified, the system must be configured so at least one of the drives that are allocated to hold quorum will remain online when the canister goes offline.

User response: Assign one or more of the drives in the control enclosure as a quorum drive. After you have configured the quorum drives, test for dependencies.

CMMVC6620E The command cannot be initiated because the drives that you have specified are in different I/O groups.

Explanation: All of the specified drives that comprise the array must be in the same I/O group.

User response: Specify one or more drives in the same I/O group, and resubmit the command.

CMMVC6621E The command cannot be initiated because the array member that you have specified already exists.

Explanation: A drive has already been configured for the specified array member. You can use the **lsarraymember** command to display the available members of an array.

User response: Specify an array member without a corresponding drive, and resubmit the command.

CMMVC6622E The command cannot be initiated because the drive has failed validation tests.

Explanation: When a drive is made a candidate, the new drive is validated to ensure that adding it to the

configuration will not adversely affect the existing or future array status. Either the current status of the drive has not allowed the validation to be performed, or the validation has failed.

User response: Fix any errors associated with the specified drive, or specify a different drive, and resubmit the command.

CMMVC6623E The command cannot be initiated because the drive validation test has timed out.

Explanation: When a drive is made a candidate, the new drive is validated to ensure that adding it to the configuration will not adversely affect the existing or future array status. The test timed out, which caused the validation to fail.

User response: Fix any errors that are associated with the specified drive, or specify a different drive, and resubmit the command.

CMMVC6624E The command cannot be initiated because the drive is not in an appropriate state to perform the task.

Explanation: The drive that you have specified is offline. A format task is permitted to an offline drive only if the drive has indicated that a format is required and connectivity to the drive is available.

User response: Fix any errors that are associated with the specified drive, or specify a different drive, and resubmit the command.

CMMVC6625E The command cannot be initiated because a task is in progress on the drive.

Explanation: A drive can complete only one task at a time. A previous task remains uncompleted. You can monitor the progress of the task using the `svcinfo lsdriveprogress` command.

User response: Wait for the previous task to complete, and resubmit the command.

CMMVC6626E The task was not initiated because a command was rejected by the drive that you have specified.

Explanation: When attempting to initiate a task, a sequence of commands is sent to the drive. One or more of these commands was rejected by the drive that you have specified.

User response: Fix any errors that are associated with the enclosure and cabling, and resubmit the command.

CMMVC6627E The enclosure that you have specified cannot be changed to managed mode because of a SAS configuration problem that is described in the event log.

Explanation: The status of the enclosure that you have specified will not allow the enclosure to be managed by the cluster.

User response: Ensure that the enclosure is online and cabled correctly, and resubmit the command.

CMMVC6628E The enclosure that you have specified cannot be changed to unmanaged mode because one or more drives are in use.

Explanation: The status of the enclosure that you have specified will not allow the enclosure to be unmanaged by the cluster.

User response: Stop using the drives, and resubmit the command.

CMMVC6630E A drive dump was not created because a command was rejected by the drive that you have specified.

Explanation: When initiating a drive dump, a sequence of commands is sent to the drive. One or more of these commands was rejected by the drive that you have specified.

User response: Fix any errors associated with the drive, enclosure, and cabling, or specify a different drive, and resubmit the command.

CMMVC6631E The task was not completed because the drive that you have specified was unavailable.

Explanation: The drive that you have specified did not have the required connectivity to complete the task.

User response: Fix any errors associated with the drive, or specify a different drive, and resubmit the command.

CMMVC6988E The command cannot be initiated because the maximum number of iSCSI qualified names (IQNs) for the cluster has been reached.

Explanation: The specified cluster is already configured with the maximum number of IQNs.

User response: None.

CMMVC6998E The command cannot be initiated because the maximum number of iSCSI qualified names (IQNs) for one or more I/O groups has been reached.

Explanation: One or more I/O groups are already

CMMVC6999E • CMMVC7011E

configured with the maximum number of IQNs.

User response: None.

ensure that all drives within the enclosure are online before resubmitting the command.

CMMVC6999E The command cannot be initiated because the maximum number of iSCSI qualified names (IQNs) for the host has been reached.

Explanation: The specified host is already configured with the maximum number of IQNs.

User response: None.

CMMVC7003E The command cannot be initiated because the power supply unit (PSU) that you have specified is offline.

Explanation: The power supply unit (PSU) that you specify must be online when you submit the command.

User response: Fix any errors associated with the specified PSU. Ensure that the PSU is online, and resubmit the command.

CMMVC7005E The command cannot be initiated because enclosures do not exist for the I/O group that you have specified.

Explanation: You have submitted a command and specified an I/O group that is not associated with an enclosure. You can submit the `svcinfo lsenclosure` command to show all of the existing enclosures and their associated I/O groups.

User response: Specify an I/O group that is associated with an enclosure when you submit the command.

CMMVC7010E The command cannot be initiated because the MDisk mode is set to Array.

Explanation: This command requires the selected MDisk to be a SAN MDisk (an MDisk that is not an array made from local drives). The selected MDisk has its mode set to Array.

User response: Use `svcinfo lsmdisk` to list the MDisks, and resubmit the command against an MDisk with a mode other than Array.

CMMVC7011E The array cannot be created because no quorum disks are currently configured.

Explanation: When creating an array, quorum disks are required to back up metadata for the array. Creating an array while no quorum disks are configured is not permitted. Quorum disks can be assigned to drives in the control enclosure automatically, or manually by using the `chquorum` command.

User response: Manage the control enclosure, and

Appendix. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features

This list includes the major accessibility features in the management GUI:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. The following screen reader has been tested: JAWS 11.
- Most of the GUI features are accessible by using the keyboard. For those features that are not accessible, equivalent function is available by using the command-line interface (CLI).
- When setting or changing an IP address on the SAN Volume Controller front panel, you can disable the fast increase function to reduce the address scrolling speed of the up and down buttons to two seconds. This feature is documented in the topic that discusses initiating cluster creation from the front panel, which is located in the IBM System Storage SAN Volume Controller Information Center and the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

Navigating by keyboard

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. You can navigate the management GUI and help system from the keyboard by using the following key combinations:

- To navigate between different GUI panels, select the Low-graphics mode option on the GUI login panel. You can use this option to navigate to all the panels without manually typing the web addresses.
- To go to the next frame, press Ctrl+Tab.
- To move to the previous frame, press Shift+Ctrl+Tab.
- To navigate to the next link, button, or topic within a panel, press Tab inside a frame (page).
- To move to the previous link, button, or topic within a panel, press Shift+Tab.
- To select GUI objects, press Enter.
- To print the current page or active frame, press Ctrl+P.
- To expand a tree node, press the Right Arrow key. To collapse a tree node, press the Left Arrow key.
- To scroll all the way up, press Home; to scroll all the way down, press End.
- To go back, press Alt+Left Arrow key.
- To go forward, press Alt+Right Arrow key.
- For actions menus:
 - Press Tab to navigate to the grid header.
 - Press the Left or Right Arrow keys to reach the drop-down field.
 - Press Enter to open the drop-down menu.
 - Press the Up or Down Arrow keys to select the menu items.
 - Press Enter to launch the action.
- For filter panes:
 - Press Tab to navigate to the filter panes.
 - Press the Up or Down Arrow keys to change the filter or navigation for nonselection.
 - Press Tab to navigate to the magnifying glass icon in the filter pane and press Enter.
 - Type the filter text.

- | – Press Tab to navigate to the red X icon and press Enter to reset the filter.
- | • For information areas:
 - | – Press Tab to navigate to information areas.
 - | – Press Tab to navigate to the fields that are available for editing.
 - | – Type your edit and press Enter to issue the change command.

Accessing the publications

- | You can find the HTML version of the IBM System Storage SAN Volume Controller information at the
- | following website:

`publib.boulder.ibm.com/infocenter/svc/ic/index.jsp`

- | You can access this information using screen-reader software and a digital speech synthesizer to hear
- | what is displayed on the screen. The information was tested using the following screen reader: JAWS
- | Version 10 or later.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

*IBM Corporation
Almaden Research
650 Harry Road
Bldg 80, D3-304, Department 277
San Jose, CA 95120-6099
U.S.A.*

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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 may 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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Adobe and the Adobe logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel logo, Intel Xeon, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Index

Special characters

-filtervalue argument 251

A

about this document
 sending comments xxi
accessibility
 keyboard 501
 repeat rate of up and down
 buttons 501
 shortcut keys 501
addhostiogr command 143
addhostport command 143
adding
 nodes 14
addmdisk command 179
addnode command 108
addvdiskcopy command 151
applydrivesoftware command 233
applydisksoftware command 187
applysoftware command 137, 379
array commands
 charray 383
 charraymember 383
 lsarray 385
 lsarrayinitprogress 388
 lsarraylba 389
 lsarraymember 390
 lsarraymembergoals 392
 lsarraymemberprogress 393
 lsarraysyncprogress 395
 mkarray 396
 overview 383
 recoverarray 397
 recoverarraybycluster 397
 rarray 398
audit log commands
 catauditlog 77
 dumpauditlog 78
 lsauditlogdumps 81
 overview 77
authentication
 SSH logins 1

B

backup and restore commands 133
backup command 133
backup commands
 backup 133
 clear 134
 help 134

C

catauditlog command 77
caterrlog command 263
caterrlogbyseqnum command 263

cfgportip command 110
changing
 passwords 71
charray command 383
charraymember command 383
chauthservice command 81
chcluster command 111
chclusterip command 114
chcontroller commands 419
chcurrentuser command 83
chdrive command 233
chemail command 92
chemailserver command 93
chemailuser command 94
chenclosure command 239
chenclosurecanister command 240
chenclosureslot command 240
chenclosureevpd command 407
cherrstate command 138
chfcconsistgrp command 193
chfcmap command 193
chhost command 144
chiogr command 115
chlicense command 375
chmdisk command 187
chmdiskgrp command 180
chnode command 118
chnodeled command 399
chpartnership command 207
chquorum command 188
chrconsistgrp command 208
chrrelationship command 208
chsnmpserver command 95
chsyslogserver command 96
chuser command 84
chusergrp command 85
chvdisk command 155
chwwnn command 409
clear command 134
cleardumps command 119, 379
clearerrlog command 139
CLI (command-line interface)
 configuring PuTTY 3
 getting started 11
 preparing SSH client on AIX or
 Linux 6
 preparing SSH client on Windows 1
 using to update cluster license 12
CLI commands
 setlocale 72
 svcinfo lscluster
 changing cluster gateway
 address 62
 changing relationship
 bandwidth 62
 displaying cluster properties 12
 modifying cluster IP address 61
 svcinfo lscurrentuser 67
 svcinfo lsfcconsistgrp 32, 34
 svcinfo lsfcmap 30, 32
 svcinfo lslicense 12

CLI commands (*continued*)

 svcinfo lsuser 67
 svcinfo lsusergrp 66
 svcinfo lsvdisk 30
 svctask chcluster
 changing cluster gateway
 address 62
 changing relationship
 bandwidth 62
 modifying cluster IP address 61
 svctask chcurrentuser 67
 svctask chfcmap 32
 svctask chlicense 12
 svctask chuser 67
 svctask chusergrp 66
 svctask mkfcconsistgrp 32
 svctask mkfcmap 67
 svctask mkuser 67
 svctask mkusergrp 66
 svctask prestartfcconsistgrp 34
 svctask rmuser 67
 svctask rmusergrp 66
 svctask startfcconsistgrp 34
cluster
 authentication
 configuring cluster iSCSI 65
 configuring for iSCSI 62
 configuring iSCSI alias 64
 configuring iSCSI authentication 65
 modifying iSCSI alias 64
 cluster commands
 addnode 108
 cfgportip 110
 chcluster 111
 chclusterip 114
 chiogr 115
 chnode 118
 cleardumps 119
 cpdumps 120
 detectmdisk 122
 mkcluster 413
 ping 123
 rmnode 123
 rmportip 125
 setclustertime 126
 setpwdreset 127
 settimezone 127
 startstats 128
 stopcluster 130
 stopstats 133
 cluster diagnostic and service-aid
 commands
 applysoftware 137
 cherrstate 138
 clearerrlog 139
 dumperrlog 139
 finderr 140
 overview 137
 setlocale 140
 svqueryclock 141
 writesernum 142

- cluster error log
 - displaying 263
- clusters
 - deleting nodes 58
 - error logs 72
 - gateway address
 - changing 62
 - logs 72
 - properties 12
 - recovering nodes 48
 - removing nodes 58
 - shutting down 72
 - updating
 - license 12
 - viewing
 - license 12
 - viewing feature logs 72
- command 237
- command-line interface (CLI)
 - configuration 2
 - configuring PuTTY 3
 - getting started 11
 - preparing SSH clients on AIX or Linux 6
 - preparing SSH clients on Windows 1
 - using to update cluster license 12
 - using to view cluster license 12
- commands
 - addhostiogrp 143
 - addhostport 143
 - addmdisk 179
 - addnode 108
 - addvdiskcopy 151
 - applydrivesoftware 233
 - applymdisksoftware 187
 - applysoftware 137, 379
 - backup 133
 - catauditlog 77
 - caterrlog 263
 - caterrlogbyseqnum 263
 - cfgportip 110
 - charray 383
 - charraymember 383
 - chauthservice 81
 - chcluster 111
 - chclustererip 114
 - chcontroller 419
 - chcurrentuser 83
 - chdrive 233
 - chemail 92
 - chemailserver 93
 - chemailuser 94
 - chenclosure 239
 - chenclosurecanister 240
 - chenclosureslot 240
 - chenclosurevpd 407
 - cherrstate 138
 - chfcconsistgrp 193
 - chfcmap 193
 - chhost 144
 - chiogrp 115
 - chlicense 375
 - chmdisk 187
 - chmdiskgrp 180
 - chnode 118
 - chnodeled 399
 - chpartnership 207
- commands (*continued*)
 - chquorum 188
 - chrconsistgrp 208
 - chrrelationship 208
 - chsnmpserver 95
 - chsyslogserver 96
 - chuser 84
 - chusergrp 85
 - chvdisk 155
 - chwwnn 409
 - clear 134
 - cleardumps 119, 379
 - clearerrlog 139
 - cpdumps 120
 - cpfiles 410
 - detectmdisk 122
 - dumpallmdiskbadblocks 189
 - dumpauditlog 78
 - dumperrlog 139, 379
 - dumpinternallog 377
 - dumpmdiskbadblocks 190
 - exit 379
 - expandvdisksize 158
 - finderr 140
 - help 134
 - includemdisk 191
 - installsoftware 411
 - leavecluster 411
 - ls2145dumps 263, 381
 - lsarray 385
 - lsarrayinitprogress 388
 - lsarraylba 389
 - lsarraymember 390
 - lsarraymembergoals 392
 - lsarraymemberprogress 393
 - lsarraysyncprogress 395
 - lsauditlogdumps 81
 - lscimomdumps 263, 381
 - lscluster 264
 - lsclustercandidate 268
 - lsclustererip 269
 - lsclustervpd 381
 - lscmdstatus 399
 - lscontroller 271
 - lscontrollerdependentvdisks 273
 - lscopystatus 263
 - lscurrentuser 274
 - lsdependentvdisks 370
 - lsdiscoverystatus 275
 - lsdrive 234
 - lsdrivelba 236
 - lsdumps 276
 - lsemailserver 277
 - lsemailuser 278
 - lsclosure 242
 - lsclosurebattery 243
 - lsclosurecanister 245
 - lsclosurepsu 247
 - lsclosureslot 248
 - lserrlogbyfconsistgrp 279
 - lserrlogbyfcmap 279
 - lserrlogbyhost 279
 - lserrlogbyiogrp 279
 - lserrlogbymdisk 279
 - lserrlogbymdiskgrp 279
 - lserrlogbynode 279
 - lserrlogbyrconsistgrp 280
- commands (*continued*)
 - lserrlogbyrrelationship 280
 - lserrlogbyvdisk 280
 - lserrlogdumps 280, 381
 - lseventlog 280
 - lsfabric 284
 - lsfconsistgrp 285
 - lsfcmap 287
 - lsfcmapcandidate 290
 - lsfcmapdependentmaps 291
 - lsfcmapprogress 290
 - lsfeaturedumps 292, 381
 - lsfiles 400
 - lsfreeextents 292
 - lshbaportcandidate 293
 - lshost 294
 - lshostiogrp 296
 - lshostvdiskmap 297
 - lsiogrp 298
 - lsiogrpcandidate 301
 - lsiogrpghost 300
 - lsiostatsdumps 302, 381
 - lsiotracedumps 302, 381
 - lsiscsiauth 302
 - lslicense 303
 - lsmdisk 304
 - lsmdiskcandidate 310
 - lsmdiskdumps 309, 382
 - lsmdiskextent 311
 - lsmdiskgrp 313
 - lsmdisklba 309
 - lsmdiskmember 317
 - lsmigrate 318
 - lsnode 319
 - lsnodecandidate 323
 - lsnodedependentvdisks 324
 - lsnodevpd 324, 382
 - lsportip 328
 - lsquorum 330
 - lsrconsistgrp 331
 - lsrrelationship 334
 - lsrrelationshipcandidate 336
 - lsrrelationshipprogress 337
 - lsrepairsevdiskcopyprogress 338
 - lsrepairvdiskcopyprogress 340
 - lsrmvdiskdependentmaps 341
 - lsroute 342
 - lsservicenodes 400
 - lsservicerecommendation 402
 - lsservicestatus 402
 - lssevdiskcopy 342
 - lssnmpserver 346
 - lssoftwaredumps 347, 382
 - lssoftwareupgradestatus 347
 - lssyslogserver 347
 - lstimezones 349
 - lsuser 349
 - lsusergp 351
 - lsvdisk 352
 - lsvdiskcopy 357
 - lsvdiskdependentmaps 361
 - lsvdiskextent 361
 - lsvdiskfcmapcopies 363
 - lsvdiskfcmapmappings 364
 - lsvdiskhostmap 364
 - lsvdisklba 365
 - lsvdiskmember 366

commands (*continued*)

- lsvdiskprogress 368
- lsvdisksyncprogress 369
- metadata 412
- migrateexts 225
- migratetoimage 226
- migratevdisk 227
- mkarray 396
- mkcluster 413
- mkemailserver 97
- mkemailuser 98
- mkfcconsistgrp 195
- mkfcmap 196
- mkhost 146
- mkmdiskgrp 181
- mkpartnership 210
- mkrconsistgrp 211
- mkrrelationship 212
- mksnmpserver 100
- mksyslogserver 101
- mkuser 86
- mkusergrp 87
- mkvdisk 160
- mkvdiskhostmap 167
- ping 123
- prestartfcconsistgrp 198, 201
- prestartfcmap 199
- recoverarray 397
- recoverarraybycluster 397
- recovervdisk 168
- recovervdiskbycluster 169
- recovervdiskbyiogrp 169
- repairsevdiskcopy 170
- repairvdiskcopy 171
- rescuenode 414
- resetleds 373
- resetpassword 414
- restore 135
- rmarray 398
- rmemailserver 102
- rmemailuser 102
- rmfcconsistgrp 200
- rmfcmap 201
- rmhost 147
- rmhostiogrp 148
- rmhostport 149
- rmmdisk 183
- rmmdiskgrp 184
- rmnode 123
- rmpartnership 214
- rmportip 125
- rmrconsistgrp 215
- rmrrelationship 215
- rmsnmpserver 103
- rmsyslogserver 103
- rmuser 88
- rmusergrp 89
- rmvdisk 172
- rmvdiskcopy 173
- rmvdiskhostmap 174
- sendinventoryemail 104
- service information 399
- service task 407
- setclustertime 126
- setdisktrace 229
- setlocale 140
- setpacedccu 415

commands (*continued*)

- setpwdreset 127
- setquorum 192
- settempsshkey 415
- settimezone 127
- settrace 230
- showtimezone 372
- shrinkvdisksize 175
- snap 414
- splitvdiskcopy 177
- startemail 104
- startfcconsistgrp 201
- startfcmap 202
- starttrconsistgrp 216
- starttrrelationship 218
- startservice 408, 416
- startstats 128
- starttrace 232
- stopcluster 130
- stopemail 105
- stopfcconsistgrp 203
- stopfcmap 204
- stopnode 416
- stoprconsistgrp 220
- stoprrelationship 221
- stopservice 417
- stopstats 133
- stoptrace 232
- svqueryclock 141
- switchrconsistgrp 222
- switchrrelationship 223
- t3recovery 417
- testemail 105
- triggerenclosureddump 250
- triggermddiskdump 193
- writeserum 142
- comments, sending xxi
- communications
 - determining between hosts and virtual disks 41
- configuring
 - iSNS server address 64
 - PuTTY 3
 - remote authentication service using CLI 65
- Connecting to the CLI using OpenSSH 7
- consistency group
 - deleting FlashCopy 35
 - stoppingFlashCopy 35
- consistency groups, Global Mirror
 - creating 38
 - deleting 39
 - modifying 38
 - starting and stopping 38
- consistency groups, Metro Mirror
 - creating 38
 - deleting 39
 - modifying 38
 - starting and stopping 38
- controller commands
 - chcontroller 419
 - overview 419
- controllers
 - changing 419
 - command 271, 419
- cpdumps command 120
- cpfiles command 410

creating

- virtual disk-to-host mappings 30
- creating users 7
- current time zone 372

D

- data migration progress
 - viewing 318
- deleting
 - nodes 58
- dependent maps
 - viewing 291
- detectmddisk command 122
- determining
 - communications between hosts and virtual disks 41
- diagnostic and service-aid commands
 - clearerrlog
 - cluster 139
 - cluster 137
 - applysoftware 137
 - cherrstate 138
 - setlocale 140
 - svqueryclock 141
 - writeserum 142
 - dumpperrlog
 - cluster 139
 - finderr
 - cluster 140
 - overview 137
- discovering
 - managed disks 18
- disks
 - migrating 55
 - migrating image mode 58
- documentation
 - improvement xxi
- drive commands
 - applydrivesoftware 233
 - chdrive 233
 - lsdrive 234
 - lsdrivelba 236
 - overview 233
- dump files
 - listing 263, 381
 - lsfeaturedumps 292
- dumpallmddiskbadblocks command 189
- dumpauditlog command 78
- dumpperrlog command 139, 379
- dumpinternallog command 377
- dumpmddiskbadblocks command 190

E

- e-mail
 - inventory reports 70
 - setting up event notification 70
- e-mail and event notification commands
 - chemailserver 93
 - chsnmpserver 95
 - chsyslogserver 96
 - mksnmpserver 100
 - mksyslogserver 101
 - rmemailserver 102
 - rmsnmpserver 103

- e-mail and event notification commands (*continued*)
 - rmsyslogserver 103
- e-mail commands
 - chemail 92
 - chemailuser 94
 - lsemailer 278
 - mkemailuser 98
 - rmemailuser 102
 - sendinventoryemail 104
 - startemail 104
 - stopemail 105
 - testemail 105
- e-mail servers
 - setting up
 - CLI 71
- email and event notification commands
 - mkemailserver 97
- email commands
 - overview 91
- enclosure commands
 - lsclosure 242
 - lsclosurecanister 245
 - lsclosurepsu 247
 - lsclosureslot 248
 - overview 239
- error log dump files
 - viewing 381
- error notification
 - SYSLOG 69
- event notification commands
 - overview 91
- exit command 379
- expanding
 - virtual disks 51, 52
- expandvdisksize command 158
- extent allocation
 - viewing 311
- extents
 - migrating
 - using the CLI (command-line interface) 53

F

- featurization settings 303
- filtering
 - FlashCopy
 - consistency groups 285
 - mappings 287, 341, 361, 363
- finderr command 140
- FlashCopy
 - consistency group
 - deleting using CLI 35
 - stopping using CLI 35
 - consistency groups
 - creating using CLI 32
 - preparing using the CLI 34
 - starting using the CLI 34
 - deleting consistency group 35
 - deleting mapping 32
 - mapping
 - deleting using CLI 32
 - stopping 32
 - mappings
 - adding to consistency group 32
 - creating using CLI 30

- FlashCopy (*continued*)
 - memory 23
 - stopping consistency group 35
- FlashCopy commands
 - chfcconsistgrp 193
 - chfcmap 193
 - mkfcconsistgrp 195
 - mkfcmap 196
 - overview 193
 - prestartfcconsistgrp 198, 201
 - prestartfcmap 199
 - rmfcconsistgrp 200
 - rmfcmap 201
 - startfcconsistgrp 201
 - startfcmap 202
 - stopfcconsistgrp 203
 - stopfcmap 204
- FlashCopy progress 290
- free extents 292
- front panel
 - password 13

G

- gateway address
 - changing 62
- Generating an SSH key pair using OpenSSH 6
- getting started
 - using the CLI (command-line interface) 11
 - using the command-line interface (CLI) 11
- Global Mirror
 - memory 23
- Global Mirror commands
 - chpartnership 207
 - chrconsistgrp 208
 - chrrelationship 208
 - mkpartnership 210
 - mkrconsistgrp 211
 - mkrrelationship 212
 - overview 207
 - rmpartnership 214
 - rmrconsistgrp 215
 - rmrrelationship 215
 - starttrconsistgrp 216
 - starttrrelationship 218
 - stoprconsistgrp 220
 - stoprrelationship 221
 - switchchrconsistgrp 222
 - switchchrrelationship 223

H

- help command 134
- host bus adaptor, unconfigured
 - viewing 293
- host commands
 - addhostiogr 143
 - addhostport 143
 - chhost 144
 - mkhost 146
 - overview 143
 - rmhost 147
 - rmhostiogr 148

- host commands (*continued*)
 - rmhostport 149
- host I/O group 296
- host objects
 - configuring using CLI 28
- hosts
 - commands 143
 - determining VDisk names 41
 - mapping virtual disks (VDisks) 30
 - viewing 294

I

- image mode VDisks
 - converting to managed mode
 - using CLI (command-line interface) 57
- includemdisk command 191
- information
 - center xix
- information commands 341, 363
 - caterrlog 263
 - caterrlogbyseqnum 263
 - chenclosure 239
 - chenclosurecanister 240
 - chenclosureslot 240
 - ls2145dumps 263
 - lscimomdumps 263
 - lscluster 264
 - lsclustercandidate 268
 - lsclusterrip 269
 - lscontroller 271
 - lscopystatus 263
 - lscurrentuser 274
 - lsdependentvdisks 370
 - lsdiscoverystatus 275
 - lsdumps 276
 - lsemailer 277
 - lsclosurebattery 243
 - lserrlogbyfcconsistgrp 279
 - lserrlogbyfcmap 279
 - lserrlogbyhost 279
 - lserrlogbyiogr 279
 - lserrlogbymdisk 279
 - lserrlogbymdiskgp 279
 - lserrlogbynode 279
 - lserrlogbyrconsistgrp 280
 - lserrlogbyrrelationship 280
 - lserrlogbyvdisk 280
 - lserrlogdumps 280
 - lseventlog 280
 - lsfabric 284
 - lsfcconsistgrp 285
 - lsfcmap 287
 - lsfcmapcandidate 290
 - lsfcmapdependentmaps 291
 - lsfcmapprogress 290
 - lsfeaturedumps 292
 - lsfreeextents 292
 - lshbaportcandidate 293
 - lshost 294
 - lshostiogr 296
 - lshostvdiskmap 297
 - lsiogr 298
 - lsiogrcandidate 301
 - lsiogrphost 300
 - lsiostatsdumps 302

information commands (*continued*)

- lsiotracedumps 302
- lsiscsiauth 302
- lslicense 303
- lsmdisk 304
- lsmdiskcandidate 310
- lsmdiskdumps 309
- lsmdiskextent 311
- lsmdiskgrp 313
- lsmdisklba 309
- lsmdiskmember 317
- lsmigrate 318
- lsnode 319
- lsnodecandidate 323
- lsnodedependentvdisks 324
- lsnodevpd 324
- lsportip 328
- lsquorum 330
- lsrconsistgrp 331
- lsrrelationship 334
- lsrrelationshipcandidate 336
- lsrrelationshipprogress 337
- lsroute 342
- lssnmpserver 346
- lssoftwaredumps 347
- lssyslogserver 347
- lstimezones 349
- lsuser 349
- lsusergrp 351
- lsvdisk 352
- lsvdiskdependentmaps 361
- lsvdiskextent 361
- lsvdiskfcmappings 364
- lsvdiskhostmap 364
- lsvdisklba 365
- lsvdiskmember 366
- lsvdiskprogress 368
 - overview 263
 - resetleds 373
 - showtimezone 372
 - triggerenclosuredump 250
- installsoftware command 411
- inventory commands
 - chcluster 111
 - chemail 92
 - mkemailuser 98
 - rmemailuser 102
 - sendinventoryemail 104
 - startemail 104
 - stopemail 105
 - testemail 105
- IP addresses
 - changing 61
- iSCSI alias
 - configuring 64
 - modifying 64
- iSNS server address
 - configuring 64

K

keyboard 501

L

- language
 - changing locale 72
- legal notices 503
- license
 - changing settings 375
 - updating
 - using the CLI (command-line interface) 12
 - viewing 303
- licensing commands
 - chlicense 375
 - dumpinternallog 377
- list dump command 261
- locale
 - changing 72
- ls2145dumps command 263, 381
- lsarray command 385
- lsarrayinitprogress command 388
- lsarraylba command 389
- lsarraymember command 390
- lsarraymembergoals command 392
- lsarraymemberprogress command 393
- lsarraysyncprogress command 395
- lsauditlogdumps command 81
- lscimomdumps command 263, 381
- lscluster command 264
- lsclustercandidate command 268
- lsclusterip command 269
- lsclustervpd command 381
- lscmdstatus command 399
- lscontroller command 271
- lscontrollerdependentvdisks command 273
- lscopystatus command 263
- lscurrentuser command 274
- lsdependentvdisks command 370
- lsdiscoverystatus command 275
- lsdrive command 234
- lsdrive1ba command 236
- lsdriveprogress 237
- lsdriveprogress command 237
- lsdumps command 276
- lsemailer command 277
- lsemaileruser command 278
- lsclosure command 242
- lsclosurebattery command 243
- lsclosurecanister command 245
- lsclosurepsu command 247
- lsclosureslot command 248
- lserrlogbyfconsistgrp command 279
- lserrlogbyfcmmap command 279
- lserrlogbyhost command 279
- lserrlogbyiogrp command 279
- lserrlogbymdisk command 279
- lserrlogbymdiskgp command 279
- lserrlogbynnode command 279
- lserrlogbyrconsistgrp command 280
- lserrlogbyrrelationship command 280
- lserrlogbyvdisk command 280
- lserrlogdumps command 280, 381
- lseventlog command 280
- lsfabric command 284
- lsfconsistgrp command 285
- lsfcmmap command 287
- lsfcmmapcandidate command 290
- lsfcmmapdependentmaps command 291
- lsfcmmapprogress command 290
- lsfeaturedumps command 292
- lsfeaturedumps commands 381
- lsfiles command 400
- lsfreeextents command 292
- lshbaportcandidate command 293
- lshost command 294
- lshostiogrp command 296
- lshostvdiskmap command 297
- lsiogrp command 298
- lsiogrpcommand candidate command 301
- lsiogrpghost command 300
- lsiostatsdumps command 302, 381
- lsiotracedumps command 302
- lsiotracedumps commands 381
- lsiscsiauth command 302
- lslicense command 303
- lsmdisk command 304
- lsmdiskcandidate command 310
- lsmdiskdumps command 309, 382
- lsmdiskextent command 311
- lsmdiskgrp command 313
- lsmdisklba command 309
- lsmdiskmember command 317
- lsmigrate command 318
- lsnode command 319
- lsnodecandidate command 323
- lsnodedependentvdisks command 324
- lsnodevpd command 324, 382
- lsportip command 328
- lsquorum command 330
- lsrconsistgrp command 331
- lsrrelationship command 334
- lsrrelationshipcandidate command 336
- lsrrelationshipprogress command 337
- lsrepairsevdiskcopyprogress command 338
- lsrepairvdiskcopyprogress command 340
- lsrmvdiskdependentmaps command 341
- lsroute command 342
- lsservicenodes command 400
- lsservicerecommendation command 402
- lsservicestatus command 402
- lssevdiskcopy command 342
- lssnmpserver command 346
- lssoftwaredumps command 347, 382
- lssoftwareupgradestatus command 347
- lssyslogserver command 347
- lstimezones command 349
- lsuser command 349
- lsusergrp command 351
- lsvdisk command 352
- lsvdiskcopy command 357
- lsvdiskdependentmaps command 361
- lsvdiskextent command 361
- lsvdiskfcmmapcopies command 363
- lsvdiskfcmappings command 364
- lsvdiskhostmap command 364
- lsvdisklba command 365
- lsvdiskmember command 366
- lsvdiskprogress command 368
- lsvdisksyncprogress command 369

M

- maintaining
 - passwords 13
- managed disk (MDisk) groups
 - creating using the CLI 19
- managed disk commands
 - applymdisksoftware 187
 - chmdisk 187
 - chquorum 188
 - includemdisk 191
 - lsquorum 330
 - overview 187
 - setquorum 192
 - triggermdiskdump 193
- managed disk group commands
 - addmdisk 179
 - chmdiskgrp 180
 - mkmdiskgrp 181
 - overview 179
 - rmmdisk 183
 - rmmdiskgrp 184
- managed disks
 - viewing disks 304, 309
 - viewing groups 313
- managed disks (MDisks)
 - adding 21
 - discovering 18
 - rebalancing access 18
 - virtual disks (VDisks)
 - relationships 42
- managed mode virtual disks
 - converting from image mode
 - using the CLI (command-line interface) 57
- mapping
 - deleting FlashCopy 32
- master console
 - configuration 2
- MDisk commands
 - dumpallmdiskbadblocks 189
 - dumpmdiskbadblocks 190
- MDisks (managed disks)
 - adding 21
 - VDisk (virtual disks) relationships 42
- MDisks *See* managed disks 179, 187
- metadata command 412
- Metro Mirror
 - memory 23
- Metro Mirror commands
 - chpartnership 207
 - chrconsistgrp 208
 - chrrelationship 208
 - mkpartnership 210
 - mkrconsistgrp 211
 - mkrcrelationship 212
 - overview 207
 - rmpartnership 214
 - rmrconsistgrp 215
 - rmrrelationship 215
 - startrcconsistgrp 216
 - startrcrelationship 218
 - stoprcconsistgrp 220
 - stoprcrelationship 221
 - switchchrconsistgrp 222
 - switchchrrelationship 223
- migrateexts command 225
- migratetoimage command 226

- migratevdisk command 227
- migrating
 - VDisks (virtual disks) 45
- migrating(volumes)
 - extents
 - using the CLI (command-line interface) 53
- migration 225
- migration commands
 - migrateexts 225
 - migratetoimage 226
 - migratevdisk 227
 - overview 225
- mkarray command 396
- mkcluster command 413
- mkemailserver command 97
- mkemailuser command 98
- mkfcconsistgrp command 195
- mkfemap command 196
- mkhost commands 146
- mkmdiskgrp command 181
- mkpartnership command 210
- mkrconsistgrp command 211
- mkrcrelationship command 212
- mksnmpserver command 100
- mksyslogserver command 101
- mkuser command 86
- mkusergrp command 87
- mkvdisk commands 160
- mkvdiskhostmap command 167

N

- nodes
 - adding 14, 108
 - addnode command 108
 - changing 118
 - chnode command 118
 - deleting 58, 123
 - removing 58
 - returning to cluster 48
 - rmnode command 123
 - viewing 319
 - general details 17

O

- OpenSSH, Connecting to the CLI
 - using 7
- OpenSSH, Generating an SSH key pair
 - using 6
- overview
 - array commands 383
 - audit log commands 77
 - backup and restore commands 133
 - cluster commands 107
 - cluster diagnostic and service-aid commands 137
 - controller commands 419
 - drive commands 233
 - dumps commands 261
 - email commands 91
 - enclosure commands 239
 - event notification commands 91
 - FlashCopy commands 193
 - host commands 143

- overview (*continued*)
 - information commands 263
 - licensing commands 375
 - managed disk commands 187
 - managed disk group commands 179
 - migration commands 225
 - secure shell 1
 - service mode commands 379
 - service mode information commands 381
 - tracing commands 229
 - user management commands 81

P

- partnerships, Global Mirror
 - creating 39
 - deleting 40
 - modifying 39
 - starting and stopping 40
- partnerships, Metro Mirror
 - creating 39
 - deleting 40
 - modifying 39
 - starting and stopping 40
- passwords
 - changing 71
 - front panel 13
- ping command 123
- plink utility
 - running 3
- port IP addresses
 - configuring 62
- prestartfcconsistgrp command 198
- prestartfcmap command 199
- PuTTY
 - configuring 3
 - generating an SSH key pair 2
 - running the plink utility 3
 - scp (pscp) 9
- PuTTY session
 - configuring for the CLI 3

Q

- quorum disks
 - setting with CLI 22

R

- reader feedback, sending xxi
- rebalancing
 - managed disks (MDisks) access 18
- recoverarray command 397
- recoverarraybycluster command 397
- recovering
 - offline virtual disks (volumes)
 - using CLI 47
- recovervdisk command 168
- recovervdiskbycluster command 169
- recovervdiskbyiogrp command 169
- related information xix
- relationships, Global Mirror
 - creating 36
 - deleting 37
 - displaying 37

- relationships, Global Mirror *(continued)*
 - modifying 36
 - starting and stopping 36
 - switching 37
- relationships, Metro Mirror
 - creating 36
 - deleting 37
 - displaying 37
 - modifying 36
 - starting and stopping 36
 - switching 37
- remote authentication
 - configuring using CLI 65
- removing
 - nodes 58
- repairing
 - space-efficient VDisk 47
- repairsevdiskcopy command 170
- repairvdiskcopy command 171
- rescuenode command 414
- resetleds command 373
- resetpassword command 414
- restore command 135
 - clear 134
 - help 134
 - restore 135
- rmarray command 398
- rmemailserver command 102
- rmemailuser command 102
- rmfcconsistgrp command 200
- rmfcmap command 201
- rmhost command 147
- rmhostiogrps command 148
- rmhostport command 149
- rmmdisk command 183
- rmmdiskgrp command 184
- rmnode command 123
- rmpartnership command 214
- rmportip command 125
- rmrconsistgrp command 215
- rmrrelationship command 215
- rmsnmpserver command 103
- rmsyslogserver command 103
- rmuser command 88
- rmusergrp command 89
- rmvdisk command 172
- rmvdiskcopy command 173
- rmvdiskhostmap command 174
- running
 - PuTTY plink utility 3

S

- SAN Volume Controller
 - front panel password 13
 - properties 17
- SAN Volume Controller library
 - related publications xix
- scanning
 - fibre-channel network 18
 - rebalancing MDisk access 18
- scp
 - PuTTY application 9
- secure shell
 - PuTTY 3

- secure shell (SSH)
 - authenticating logins 1
 - client
 - AIX or Linux 6
 - Windows 1
 - creating keys 2
 - overview 1
 - secure shell client
 - preparing for CLI on AIX 6
 - preparing for CLI on Linux 6
 - security 1
 - sending
 - comments xxi
 - sendinventoryemail command 104
 - service commands
 - chenclosurevpd 407
 - chnodeled 399
 - chwwnn 409
 - installsoftware 411
 - leavecluster 411
 - lscmdstatus 399
 - lsfiles 400
 - lsservicenodes 400
 - lsservicerecommendation 402
 - lsservicestatus 402
 - metadata 412
 - rescuenode 414
 - resetpassword 414
 - setpacedccu 415
 - settempsshkey 415
 - snap 414
 - startservice 408, 416
 - stopnode 416
 - stopservice 417
 - t3recovery 417
 - service information commands 399
 - service mode
 - commands 379
 - information commands 381
 - service mode commands
 - applysoftware 379
 - cleardumps 379
 - dumperrlog 379
 - exit 379
 - overview 379
 - service mode information commands
 - ls2145dumps 381
 - lscimomdumps 381
 - lsclustervpd 381
 - lserrlogdumps 381
 - lsfeaturedumps 381
 - lsiotstatsdumps 381
 - lsiotracedumps 381
 - lsmdiskdumps 382
 - lsnodevpd 382
 - lssoftwaredumps 382
 - overview 381
 - service task commands 407
 - cpfiles 410
 - setclustertime command 126
 - setdisktrace command 229
 - setlocale command 140
 - setpacedccu command 415
 - setpwdreset command 127
 - setquorum command 192
 - settempsshkey command 415
 - settimezone command 127

- setting
 - quorum disks 22
- settings
 - e-mail server 71
 - error notification 69
 - event notification 68
- settrace command 230
- shortcut keys 501
- showtimezone command 372
- shrinkvdisksize command 53, 175
- snap command 414
- SNMP traps 68
- software
 - upgrading using the command-line interface (CLI) 73
- software packages
 - listing 382
 - viewing 347
- splitvdiskcopy command 177
- SSH (secure shell)
 - client system
 - preparing to issue CLI commands 6
- SSH *See* secure shell 1
- SSH keys
 - creating 2
- startemail command 104
- startfcconsistgrp command 201
- startfcmap command 202
- starttrconsistgrp command 216
- starttrrelationship command 218
- startservice command 408, 411, 416
- startstats command 128
- starttrace command 232
- stopcluster command 130
- stopemail command 105
- stopfcconsistgrp command 203
- stopfcmap command 204
- stopnode command 416
- stopping
 - FlashCopy mapping 32
- stoprconsistgrp command 220
- stoprrelationship command 221
- stopservice command 417
- stopstats command 133
- stoptrace command 232
- subnet mask
 - changing 62
- svqueryclock command 141
- switchrconsistgrp command 222
- switchrrelationship command 223
- SYSLOG 69
- system log
 - information 69

T

- t3recovery command 417
- testemail command 105
- time
 - setting cluster
 - using the CLI (command-line interface) 11
- time zones 349
- tracing commands
 - overview 229
 - setdisktrace 229

- tracing commands (*continued*)
 - settrace 230
 - starttrace 232
 - stoptrace 232
- trademarks 505
- triggerenclosedump command 250
- triggerdiskdump command 193

U

- Updating
 - license
 - using the CLI (command-line interface) 12
- upgrading
 - software using the command-line interface (CLI) 73
- user groups
 - creating using CLI 66
 - modifying using CLI 66
- user management commands
 - chauthservice 81
 - chcurrentuser 83
 - chuser 84
 - chusergrp 85
 - mkuser 86
 - mkusergrp 87
 - overview 81
 - rmuser 88
 - rmusergrp 89
- users
 - creating 7
 - creating using CLI 67
 - modifying using CLI 67

V

- validating
 - VDisk copies 45
- VDisk (virtual disk)
 - expanding 52
- VDisk (virtual disks)
 - determining mappings 42
- VDisk copies
 - validating 45
- VDisk Mirroring
 - memory 23
- VDisks (virtual disks)
 - adding a copy 27
 - converting
 - from image mode to managed mode 57
 - determining name of 41
 - expanding 51
 - listing node dependent 41
 - MDisks (managed disks)
 - relationships 42
 - migrating 45, 56
 - recovering 49
 - using the CLI 49
 - viewing 352
- VDisks (virtual disks) (volume)
 - creating 24
- VDisks (virtual disks) (volume)
 - deleting a copy 28

- VDisks (volumes)
 - recovering from offline
 - using CLI 47
- VDisks *See* virtual disks 151
- viewing
 - clusters 264
 - Global Mirror
 - consistency groups 331
 - relationships 334
 - I/O groups 298
 - Metro Mirror
 - consistency groups 331
 - relationships 334
- Viewing
 - license
 - using the CLI (command-line interface) 12
- virtual disk
 - viewing FlashCopy mappings 364
- virtual disk (VDisks) extent
 - viewing 361
- virtual disk commands
 - addvdiskcopy 151
 - chvdisk 155
 - expandvdisksize 158
 - lscontrollerdependentvdisks 273
 - lsrepairvdiskcopyprogress 338
 - lsrepairvdiskcopyprogress 340
 - lssevdiskcopy 342
 - lsvdiskcopy 357
 - lsvdisksyncprogress 369
 - mkvdisk 160
 - mkvdiskhostmap 167
 - recovervdisk 168
 - recovervdiskbycluster 169
 - recovervdiskbyiogrp 169
 - repairsevdiskcopy 170
 - repairvdiskcopy 171
 - rmvdisk 172
 - rmvdiskcopy 173
 - rmvdiskhostmap 174
 - shrinkvdisksize 175
 - splitvdiskcopy 177
- virtual disks
 - copying 151
 - creating 160
 - removing 173
 - viewing disks 365
- virtual disks (VDisks) 52
 - adding a copy 27
 - converting
 - from image mode to managed mode 57
 - determining mappings 42
 - determining name of 41
 - managed disks (MDisks)
 - relationships 42
 - migrating 45
 - recovering 49
 - shrinkvdisksize command 53
 - using the CLI 49
- virtual disks (VDisks)(volume)
 - deleting a copy 28
- virtual disks commands
 - overview 151
- vital product data (VPD)
 - listing 381

- vital product data (VPD) (*continued*)
 - viewing 324
- volumes (VDisks)
 - recovering from offline
 - using CLI 47

W

- writesernum command 142



Printed in USA

GC27-2287-00

