

IBM Spectrum Virtualize for SAN Volume Controller and
Storwize Family

Command-Line Interface User's Guide



**** MISSING FILE ****

This file was generated during the publishing process

This edition applies to IBM Spectrum Virtualize for SAN Volume Controller and Storwize Family Version 7.8.1, and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces GC27-2287-13.

© **Copyright IBM Corporation 2003, 2017.**

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
Accessibility	xi
Emphasis	xi
SAN Volume Controller library and related publications	xi
IBM Publications Center	xiii
Sending your comments	xiii
Syntax diagrams	xiii
Terminology	xv
CLI special characters.	xvi
Using wildcards in the CLI	xvi
Data types and value ranges	xvii
CLI commands and parameters	xxi
CLI flags	xxii
CLI messages	xxiii
CLI deprecated and discontinued commands	xxiii
Understanding capacity indicators	xxv
Attributes of the -filtervalue parameters	xxv

Chapter 1. Setting up an SSH client. 1

Setting up an SSH client on a Windows host.	2
Generating an SSH key pair using PuTTY	2
Configuring a PuTTY session for the CLI	3
Connecting to the CLI using PuTTY	4
Starting a PuTTY session for the CLI	6
Preparing the SSH client on an AIX or Linux host	6
Generating an SSH key pair using OpenSSH.	7
Connecting to the CLI using OpenSSH	8
Working with local and remote users	8
UNIX commands available in interactive SSH sessions	9
Copying the software update files by using PuTTY pscp or openssh scp	9

Chapter 2. Using the CLI 11

Setting the clustered system time by using the CLI	11
Setting cluster date and time	12
Viewing and updating license settings by using the CLI	12
Displaying clustered system properties by using the CLI	13
Maintaining passwords using the CLI	14
Using the dump commands to work with directories	15
Re-adding a repaired node to a clustered system by using the CLI.	16
Displaying node properties by using the CLI	20
Discovering MDiskS using the CLI	20
Creating storage pools using the CLI	22
Adding MDiskS to storage pools using the CLI	24
Setting a quorum disk using the CLI	25

Modifying the amount of available memory for Copy Services, Volume Mirroring, and RAID arrays by using the CLI	26
Creating volumes using the CLI	28
Adding a copy to a volume	31
Deleting a copy from a volume.	31
Configuring host objects using the CLI	31
Creating host mappings by using the CLI	33
Creating FlashCopy mappings using the CLI	33
Preparing and starting a FlashCopy mapping using the CLI.	34
Stopping FlashCopy mappings by using the CLI	35
Deleting a FlashCopy mapping using the CLI	36
Creating a FlashCopy consistency group and adding mappings using the CLI	36
Preparing and starting a FlashCopy consistency group using the CLI	37
Stopping a FlashCopy consistency group using the CLI.	39
Deleting a FlashCopy consistency group using the CLI.	39
Creating Metro Mirror, Global Mirror, or active-active relationships using the CLI.	40
Modifying Metro Mirror, Global Mirror, or active-active relationships using the CLI.	40
Starting and stopping Metro Mirror, Global Mirror, or active-active relationships using the CLI	41
Displaying the progress of Metro Mirror, Global Mirror, or active-active relationships using the CLI	42
Switching Metro Mirror or Global Mirror relationships using the CLI	42
Deleting Metro Mirror, Global Mirror, or active-active relationships using the CLI.	43
Creating Metro Mirror, Global Mirror, or active-active consistency groups using the CLI.	43
Modifying Metro Mirror, Global Mirror, or active-active consistency groups using the CLI.	43
Starting and stopping Metro Mirror, Global Mirror, or active-active consistency-group copy processes using the CLI	44
Deleting Metro Mirror, Global Mirror, or active-active consistency groups using the CLI.	44
Creating Metro Mirror and Global Mirror partnerships using the CLI	45
Modifying Metro Mirror and Global Mirror partnerships using the CLI	46
Starting and stopping Metro Mirror and Global Mirror partnerships using the CLI.	47
Deleting Metro Mirror and Global Mirror partnerships using the CLI	47
Determining the WWNNs of a node using the CLI	47
Listing node-dependent volumes using the CLI	48
Determining the volume name from the device identifier on the host	49

Determining the host that a volume maps	50
Determining the relationship between volume and MDisk using the CLI	50
Determining the relationship between MDisk and controller LUNs using the CLI	50
Increasing the size of your system by using the CLI	51
Adding a node to increase the size of the system	51
Validating and repairing mirrored volume copies by using the CLI.	54
Repairing a thin-provisioned volume using the CLI	56
Recovering offline volumes using the CLI	56
Recovering a node and returning it to the system by using the CLI	57
Recovering offline volumes using the CLI	58
Moving offline volumes to their original I/O group using the CLI	58
Recording WWPN changes of replaced host HBAs	59
Expanding volumes using the CLI.	60
Expanding a volume that is mapped to an AIX host	61
Expanding a volume that is mapped to a Microsoft Windows host by using the CLI	61
Shrinking a volume using the CLI.	62
Migrating extents using the CLI	63
Migrating volumes between pools using the CLI	64
Moving a volume between I/O groups using the CLI	66
Creating an image-mode volume using the CLI	67
Migrating data to an image mode volume using the CLI	68
Deleting a node from a system by using the CLI	68
Completing the system maintenance procedure by using the CLI.	70
Modifying system IP addresses using the CLI	71
Changing the system gateway address by using the CLI	72
Changing the relationship bandwidth for a system by using the CLI	72
Configuring the system for iSCSI hosts	73
Configuring or modifying an iSCSI alias by using the CLI.	74
Configuring the iSNS server address using the CLI	75
Configuring system iSCSI authentication using the CLI.	75
Configuring remote authentication service by using the CLI.	76
Configuring remote authentication service with Lightweight Directory Access Protocol (LDAP) using the CLI.	76
Changing user groups.	77
Changing users	78
Setting up SNMP notifications using the CLI	79
Setting up syslog notifications using the CLI	80
Setting up email event notifications and inventory reports using the CLI	80
Setting up email servers using the CLI	82
Changing user passwords using the CLI.	82
Changing the locale setting using the CLI	83
Viewing the feature log using the CLI	83
Analyzing the error log using the CLI	83

Shutting down a system by using the CLI	84
Updating the software automatically using the CLI	85

Chapter 3. Array commands 89

charray	89
charraymember	90
lsarray	93
lsarrayinitprogress.	100
lsarraylba.	101
lsarraymember	103
lsarraymembergoals	106
lsarraymemberprogress	109
lsarrayrecommendation	111
lsarraysyncprogress	114
lspotentialarraysize	115
mkarray	117
mkdistributedarray	120
recoverarray.	123
recoverarraybycluster (Discontinued)	123
recoverarraybysystem	123
rmarray	124

Chapter 4. Audit log commands 125

catauditlog	125
dumpauiditlog	126
lsauiditlogdumps (Deprecated).	128

Chapter 5. Backup and restore commands 129

svconfig.	129
backup	130
clear	131
cron	132
recover	132
restore.	133
restore.	135

Chapter 6. Cloud account commands 137

chcloudaccountawss3.	137
chcloudaccountswift	139
lscloudaccount	142
lscloudaccountusage	144
lscloudaccountimportcandidate	146
mkcloudaccountawss3	147
mkcloudaccountswift.	148
rmcloudaccount	150
testcloudaccount	150

Chapter 7. Clustered system commands 153

addnode (SAN Volume Controller only)	153
addiscsistorageport	156
cfgportip	157
chbanner	163
chcluster (Discontinued).	165
chiogrp	165
chnode	168
chnodebattery	171
chnodebootdrive	171

chnodehw (SVC) / chnodecanisterhw (Storwize family products)	173
chquorum	174
chsecurity	176
chsite	176
chsystem	177
chsystemcert	184
chsystemip	186
chthrottle	189
cleardumps	189
cpdumps	191
detectscsistorageportcandidate	192
dumpconfig (Discontinued).	194
help	194
lsclustercandidate (Discontinued).	195
lscluster (Discontinued)	195
lsclusterip (Discontinued)	195
lsclusterstats (Discontinued)	195
lsdiscoverystatus	195
lsfabric	197
lsfcportcandidate	200
lsiscsistorageport	201
lsiscsistorageportcandidate	204
lsigrp	206
lshbaportcandidate (Deprecated)	210
lsiogrphost	210
lsiogrpcandidate	210
lsiostatsdumps (Deprecated)	211
lsiotracedumps (Deprecated)	211
lsnode (SVC) / lsnodecanister (Storwize family products)	211
lsnodebattery	216
lsnodecandidate (SAN Volume Controller).	219
lsnodedependentvdisks (Deprecated)	220
lsnodehw (SVC) / lsnodecanisterhw (Storwize family products)	221
lsnodestats (SVC) / lsnodecanisterstats (Storwize family products)	223
lsnodevpd (SVC) / lsnodecanistervpd (Storwize family products)	231
lsportusb	240
lsportip	242
lsportfc	250
lsportsas	252
lsquorum	254
lsroute	256
lstimezones	257
lssasportcandidate	258
lssecurity	259
lssite	260
lsthrottle	261
lssystem	263
lssystemcert	270
lssystemip	272
lssystemstats	275
lstargetportfc	279
mkquorumapp	281
mkthrottle	282
ping	284
rmiscsistorageport	285

rmnode (SVC) / rmnodecanister (Storwize family products)	285
rmportip	288
rmthrottle	289
setclustertime (Discontinued)	289
setsystemtime	289
setpwdreset	290
settimezone	290
showtimezone	291
startstats	292
stopstats (Deprecated)	293
stopcluster (Discontinued)	293
stopsystem	293

Chapter 8. Clustered system

diagnostic and service-aid commands 297

applysoftware	297
caterrlog (Deprecated)	300
caterrlogbyseqnum (Deprecated)	300
cherrstate (Deprecated)	300
chdnsserver	300
cheventlog	301
chsyslogserver	301
clearerrlog	302
cpfabricdumps (Discontinued).	303
dumperrlog	303
finderr	304
setevent (Discontinued)	304
lscimomdumps (Deprecated)	304
lscopystatus	304
lsdumps	305
lsdnsserver	307
lserrlogbyfcconsistgrp (Deprecated)	308
lserrlogbyfcmap (Deprecated)	308
lserrlogbyhost (Deprecated)	308
lserrlogbyiogrp (Deprecated)	308
lserrlogbymdisk (Deprecated)	308
lserrlogbymdiskgrp (Deprecated)	308
lserrlogbynode (Deprecated)	308
lserrlogbyrconsistgrp (Deprecated)	308
lserrlogbyrrelationship (Deprecated)	309
lserrlogbyvdisk (Deprecated)	309
lserrlogdumps (Deprecated)	309
lsfeaturedumps (Deprecated)	309
lseventlog	309
lsservicestatus	315
lssyslogserver	320
lssoftwareumps (Deprecated).	321
lssoftwareupgradestatus (Deprecated)	321
lsupdate	322
mkdnsserver	325
mksyslogserver	325
rmdnsserver	327
rmsyslogserver	327
setlocale	328
svqueryclock	329
writesernum	329

Chapter 9. Controller command . . . 331

chcontroller	331
------------------------	-----

lscontroller	332
lscontrollerdependentvdisks	336

Chapter 10. Drive commands 339

applydrivesoftware	339
chdrive	342
lsdrive.	343
lsdriveclass	349
lsdrivelba.	351
lsdriveprogress	352
lsdriveupgradeprogress	354
triggerdrivedump	356

Chapter 11. Email and event notification commands 359

chemail	359
chemailserver	361
chemailuser	362
chsnmpserver	363
lsemailserver	364
lsemailuser	365
lssnmpserver	366
mkemailserver	367
mkemailuser	368
mksnmpserver	370
rmemailserver	371
rmemailuser	371
rmsnmpserver	372
sendinventoryemail	372
setemail (Discontinued)	373
startemail.	373
stopemail.	374
testemail	374

Chapter 12. Enclosure commands . . . 377

addcontrolenclosure	377
chenclosure	377
chenclosurecanister	378
chenclosuredisplaypanel.	380
chenclosurepsu	380
chenclosuresem.	381
chenclosureslot	381
(satask) chenclosurevpd (Deprecated)	383
lsenclosure	383
lsenclosurebattery	386
lscontrolenclosurecandidate (Storwize family products only)	389
lsenclosurecanister.	390
lsenclosurechassis (Flex V7000 Storage Node).	393
lsenclosuredisplaypanel	394
lsenclosurefanmodule	396
lsenclosurepsu	398
lsenclosuresem	400
lsenclosureslot	402
lsenclosurestats.	405
lssasfabric	409
resetleds	411
triggerenclosuredump	412

Chapter 13. Encryption commands 413

chencryption	413
lsencryption	415

Chapter 14. Licensing and featurization commands. 419

activatefeature	419
chlicense	420
deactivatefeature	422
lsfeature	423
lslicense	425

Chapter 15. FlashCopy commands 429

chfcconsistgrp	429
chfcmap	429
lsfcconsistgrp	431
lsfcmap	434
lsfcmapcandidate	438
lsfcmapprogress	438
lsfcmapdependentmaps	439
lsrmvdiskdependentmaps	440
mkfcconsistgrp	441
mkfcmap	442
prestartfcconsistgrp	444
prestartfcmap	446
rmfcconsistgrp	447
rmfcmap	448
startfcconsistgrp	448
startfcmap	450
stopfcconsistgrp	451
stopfcmap	453

Chapter 16. Host commands 455

addhostclustermember	455
addhostiogrp	456
addhostport	456
chhost	458
chhostcluster	459
lshost	461
lshostcluster.	465
lshostclustermember	467
lshostclustervolumemap.	468
lshostiogrp	470
lsiscsiauth	471
mkhost	472
mkhostcluster	474
mkvolumehostclustermmap	475
rmhost	476
rmhostcluster	477
rmhostclustermember	478
rmvolumehostclustermmap	479
rmhostiogrp	480
rmhostport	481

Chapter 17. Information commands 483

ls2145dumps (Deprecated)	483
lsconfigdumps (Discontinued).	483
lssshkeys (Discontinued)	483

Chapter 18. Key server commands 485

chkeyserver	485
chkeyserverisklm	486
lskeyserver	487
lskeyserverisklm	488
mkkeyserver	490
rmkeyserver	491
testkeyserver	492

Chapter 19. Livedump commands . . 493

cancellivedump	493
lslivedump	493
preplivedump	494
triggerlivedump	494

Chapter 20. Managed disk commands 497

addmdisk	497
applymdisksoftware (Discontinued)	498
chmdisk	498
detectmdisk	500
dumpallmdiskbadblocks	501
dumppmdiskbadblocks	502
includemdisk	503
lsmdisk	504
lsmdiskdumps (Deprecated)	510
lsmdisklba	510
lsmdiskcandidate	511
lsmdiskextent	512
lsmdiskmember	514
setquorum (Deprecated)	515
triggermdiskdump (Discontinued)	516

Chapter 21. Storage pool commands 517

chmdiskgrp	517
lsfreeextents	519
lsmdiskgrp	519
mkmdiskgrp	527
rmmdisk	531
rmmdiskgrp	532

Chapter 22. Copy services commands 535

chpartnership	535
chrconsistgrp	538
chrcrelationship	540
lspartnership	544
lspartnershipcandidate	548
lsrconsistgrp	549
lsrcrelationship	552
lsrcrelationshipcandidate	556
lsrcrelationshipprogress	557
mkfcpartnership	558
mkippartnership	559
mkpartnership (Discontinued)	561
mkrconsistgrp	561
mkrcrelationship	562
rmpartnership	566
rmrconsistgrp	566
rmrcrelationship	567
startrcconsistgrp	568
startrcrelationship	571

stoprcconsistgrp	573
stoprcrelationship	576
switchrcconsistgrp	578
switchrcrelationship	579

Chapter 23. Migration commands . . . 581

lsmigrate	581
migrateexts	582
migratetoimage	583
migratevdisk	585

Chapter 24. Service information commands 587

host	587
sainfo lsbootdrive	587
sainfolscmdstatus	589
lsfiles	590
lshardware	591
lsservicenodes	593
lsservicerecommendation	594
lsservicestatus	595
traceroute	600

Chapter 25. Service mode commands (Discontinued) 601

applysoftware (Discontinued)	601
svcservicemodetask cleardumps (Discontinued)	601
svcservicemodetask dumperrlog (Discontinued)	601
exit (Discontinued)	601

Chapter 26. Service mode information commands (Discontinued) 603

ls2145dumps (Discontinued)	603
lscimomdumps (Discontinued)	603
lsclustervpd (Discontinued)	603
lserrlogdumps (Discontinued)	603
lsfeaturedumps (Discontinued)	603
lsiostatsdumps (Discontinued)	603
lsiotracedumps (Discontinued)	603
lsmdiskdumps (Discontinued)	603
lssoftwaredumps (Discontinued)	603

Chapter 27. Service task commands 605

satask chbootdrive	605
chnodeled	605
chserviceip	607
satask chvpd	609
chwwnn	611
cpfiles	612
dumpinternallog (Discontinued)	613
installsoftware	613
leavecluster	614
metadata	615
(satask) mkcluster	616
mkcluster (Deprecated)	617
satask overridequorum	617
rescuenode	618
resetpassword	618
restartservice	619

(satask) setlocale	620
setpacedccu	621
settempsshkey	621
snap	622
startservice	623
stopnode	623
stopservice	624
t3recovery	625

Chapter 28. User management

commands	627
chauthservice	627
chcurrentuser	629
chldap	630
chldapserver	633
chuser	634
chusergrp	636
lscurrentuser	636
lsldap	637
lsldapserver	638
lsuser	640
lsusergrp	641
mkldapserver	643
mkuser	644
mkusergrp	645
rmldapserver	648
rmuser	648
rmusergrp	649
testldapserver	649

Chapter 29. Volume commands. . . . 653

addvolume	653
addvdiskcopy	655
addvdiskaccess	661
analyzevdisk	662
analyzevdiskbysystem	663
backupvolume	664
backupvolume	665
chvdisk	666
chvolume	670
expandvdisksize	671
lsdependentvdisk	674
lshostvdiskmap	675
lsmetadatavdisk	677
lsrepairsevdiskcopyprogress	678
lsrepairvdiskcopyprogress	679
lssevdiskcopy	681
lsvdisk	687
lsvdiskaccess	701
lsvdiskanalysis	703
lsvdiskanalysisprogress	706
lsvdiskcopy	707

lsvdiskdependentmaps	713
lsvdiskextent	713
lsvdiskfcmapcopies	715
lsvdiskfcmappings	716
lsvdiskhostmap	717
lsvdisklba	718
lsvdiskmember	719
lsvdiskprogress	721
lsvdisksyncprogress	722
lsvolumebackup	723
lsvolumebackupgeneration	725
lsvolumebackupprogress	727
lsvolume	729
lsvolume	729
lsvolume	731
mkmetadatavdisk	733
mkvdisk	734
mkvdiskhostmap	743
mkvolume	745
mkvolume	748
mkimagevolume	749
movevdisk	751
recovervdisk	753
recovervdiskbycluster (Discontinued)	754
recovervdiskbyiogrp	754
recovervdiskbysystem	754
repairsevdiskcopy	755
repairvdiskcopy	756
restorevolume	757
rmvdisk	758
rmmetadatavdisk	761
rmvdiskcopy	762
rmvdiskaccess	763
rmvdiskhostmap	763
rmvolume	764
rmvolume	765
shrinkvdisksize	766
rmvolume	768
rmvolume	770
splitvdiskcopy	771

Chapter 30. Command-line interface messages 775

Appendix. Accessibility features for SAN Volume Controller 945

Notices 947

Trademarks 949

Index 951

Tables

1. IBM websites for help, services, and information	xii
2. SAN Volume Controller library	xii
3. IBM documentation and related websites	xiii
4. Syntax diagrams	xiv
5. Abbreviations	xv
6. Data types	xvii
7. Capacity indicators	xxv
8. UNIX commands for interactive SSH sessions	9
9. Maximum volume capacity by extent size	23
10. Examples of memory required	26
11. RAID level comparisons	27
12. Volume copy resynchronization rates	30
13. chararraymember combination options	92
14. Array output	95
15. lsarrayinitprogress output	101
16. lsarraylba output	102
17. lsarraymember output	104
18. lsarraymembergoals output	107
19. lsarraymemberprogress output	110
20. lsarrayrecommendation output	112
21. lsarraysyncprogress output	115
22. lspotentialarraysize output	117
23. lscloudaccount output	142
24. lscloudaccountusage output	145
25. lscloudaccountimportcandidate output	146
26. Memory required for RAID arrays, Copy Services, and volume mirroring	167
27. RAID level comparisons	167
28. Number of extents reserved by extent size	175
29. IP address list formats	188
30. lsfcportcandidate output	201
31. lsiscsistorageport output	202
32. lsiscsistorageportcandidate output	205
33. lsnode or lsnodecanister attribute values	213
34. lsnodebattery attribute values	217
35. lsnodecandidate outputs	220
36. Attribute values for lsnodehw and lsnodecanisterhw	221
37. Attribute values for lsnodestats or lsnodecanister	224
38. Stat_name field values	226
39. Attribute values for lsnodevpd and lsnodecanistervpd	232
40. lsportusb output	240
41. lsportip output	244
42. lsportfc output	251
43. lsportsas output	253
44. lsquorum output	255
45. lssasportcandidate output	259
46. lssecurity attribute values	260
47. lssite attribute values	261
48. lsthrottle output	262
49. lssystem output	263
50. lssystemcert output	271
51. lssystemip output	274
52. lssystemstats attribute values	276
53. Stat_name field values	277
54. lstargetportfc output	280
55. lsdsnserver output	307
56. lseventlog output	311
57. lsservicestatus output	315
58. lssyslogserver output	321
59. lscontroller output	334
60. lsdrive output	345
61. lsdriveclass output	350
62. lsdrivelba output	352
63. lsencllosure output	384
64. lsencllosurebattery outputs	387
65. lscontrolenclosurecandidate attribute values	389
66. lsencllosurecanister output	391
67. lsencllosurechassis outputs	394
68. ? output	395
69. lsencllosurefanmodule attribute values	397
70. lsencllosurepsu output	399
71. lsencllosuresem output	401
72. lsencllosureslot output	403
73. lsencllosurestats outputs	407
74. Stat_name field values	409
75. lssasfabric output	410
76. lsencryption output	416
77. lsfeature outputs	424
78. lslicense output	426
79. Relationship between the cleaning <i>rate</i> , data rate, and grains per second values	431
80. Relationship between the cleaning <i>rate</i> , data rate, and grains per second values	444
81. lshost output	464
82. lshostcluster output	466
83. lshostclustermember output	467
84. lshostclustervolumemap output	469
85. lskeyserver output	487
86. lskeyserverisklm output	489
87. lsivedump outputs	494
88. MDisk output	506
89. lsmdisklba command output	511
90. Parameter differences for child pools and parent pools	518
91. Parameter differences for child pools and storage pools	529
92. lspartnership attribute values	545
93. lsrrconsistgrp command output values	550
94. lsrrrelationship command attributes and values	554
95. stoprrconsistgrp consistency group states	575
96. stoprrrelationship consistency group states	577
97. lsbootdrive attribute values	588
98. lshardware attribute values	592
99. lsservicenodes outputs	593
100. lsservicestatus output	595
101. lsldap attribute values	638
102. lsldapserver attribute values	639

103. testldapservice attribute values	650	111. lsvdisklba command output scenarios	719
104. Storage pool Easy Tier settings.	658	112. lsvolumebackup output	725
105. Relationship between the <i>syncrate</i> value and the data copied per second	660	113. lsvolumebackupgeneration output	727
106. Relationship between the <i>syncrate</i> value and the data copied per second	669	114. lsvolumebackupprogress output	728
107. lsmetadataavdisk output	678	115. lsvolumegroup output	730
108. Easy Tier status values	693	116. lsvolumerestoreprogress output	732
109. lsvdiskanalysis output	704	117. Relationship between the <i>syncrate</i> value and the data copied per second	741
110. lsvdiskanalysisprogress output.	706	118. Accepted IP address formats	913

About this guide

This publication provides information that helps you configure and use the IBM Spectrum Virtualize™ for SAN Volume Controller and Storwize® Family.

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.

Accessibility

IBM® strives to provide products with usable access for everyone, regardless of age or ability.

This product uses standard Windows navigation keys.

For more information, see the accessibility features topic in the Reference section.

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.
Bold monospace	Text in bold monospace represents 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 system.
Monospace	Text in monospace identifies the data or commands that you type, samples of command output, examples of program code or messages from the system, or names of command flags, parameters, arguments, and name-value pairs.

SAN Volume Controller library and related publications

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

IBM Knowledge Center for SAN Volume Controller

The information collection in the IBM Knowledge Center contains all of the information that is required to install, configure, and manage the system. The information collection in the IBM Knowledge Center is updated between product releases to provide the most current documentation. The information collection is available at the following website:

<http://www.ibm.com/support/knowledgecenter/STPVGU>

SAN Volume Controller library

Unless otherwise noted, the publications in the library are available in Adobe portable document format (PDF) from a website.

ibm.com/shop/publications/order

Click **Search for publications** to find the online publications you are interested in, and then view or download the publication by clicking the appropriate item.

Table 1 lists websites where you can find help, services, and more information.

Table 1. IBM websites for help, services, and information

Website	Address
Directory of worldwide contacts	http://www.ibm.com/planetwide
Support for SAN Volume Controller (2145)	www.ibm.com/support
Support for IBM System Storage® and IBM TotalStorage products	www.ibm.com/support/

Each PDF publication in the Table 2 library is also available in the IBM Knowledge Center by clicking the number in the “Order number” column:

Table 2. SAN Volume Controller library

Title	Description	Order number
<i>IBM SAN Volume Controller Model 2145-SV1 Hardware Installation Guide</i>	The guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller model 2145-SV1.	GI13-4547
<i>IBM SAN Volume Controller Hardware Maintenance Guide</i>	The guide provides the instructions that the IBM service representative uses to service the SAN Volume Controller hardware, including the removal and replacement of parts.	GC27-2283
<i>IBM SAN Volume Controller Troubleshooting Guide</i>	The guide describes the features of each SAN Volume Controller model, explains how to use the front panel or service assistant GUI, and provides maintenance analysis procedures to help you diagnose and solve problems with the SAN Volume Controller.	GC27-2284
<i>IBM Spectrum Virtualize for SAN Volume Controller and Storwize Family Command-Line Interface User's Guide</i>	The guide describes the commands that you can use from the SAN Volume Controller command-line interface (CLI).	GC27-2287

IBM documentation and related websites

Table 3 on page xiii lists websites that provide publications and other information about the SAN Volume Controller or related products or technologies. The IBM Redbooks® publications provide positioning and value guidance, installation and implementation experiences, solution scenarios, and step-by-step procedures for various products.

Table 3. IBM documentation and related websites

Website	Address
IBM Publications Center	ibm.com/shop/publications/order
IBM Redbooks publications	www.redbooks.ibm.com/

Related accessibility information

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

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

IBM Publications Center

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

The IBM Publications Center website offers customized search functions to help you find the publications that you need. You can view or download publications at no charge. Access the IBM Publications Center through the following website:

ibm.com/shop/publications/order

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 Spectrum Virtualize for SAN Volume Controller and Storwize Family Command-Line Interface User's Guide*
 - Publication form number: GC27-2287-01
 - 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.

Table 4 on page xiv 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.

Table 4. Syntax diagrams

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 lefttoright, toptobottom, following the main path line.
Keyword		Represents the name of a command, flag, parameter, or argument. A keyword is not in italics. Spell a keyword exactly as it is shown in the syntax diagram.
Required keywords		Indicate the parameters or arguments 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.

Table 4. Syntax diagrams (continued)

Element	Syntax	Description
Quotation mark delimiters	<pre> >>—d— —"—ess—==—EssId—host—==—————> >—'Host Name'— —profile—==—ProfileName————> >—"—————>< </pre>	Indicates the start and end of a parameter or argument that contains multiple values. Enclose one or more namevalue pairs in a set of double quotation marks for a particular parameter or argument. If the value of a parameter or namevalue pair contains a blank or white space, enclose the entire value in a set of single quotation marks.
Equalsign operator	<pre> >>—"—ess—==—EssId— —profile—==—————> >—ProfileName—"—————>< </pre>	Separates a name from its value in a namevalue pair.
Syntax fragment	<pre> >>—Fragment Name—————>< Fragment name: —(—fragment details—)—————— </pre>	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 (CLI) operations.

Table 5 shows the abbreviations and offers a brief description.

Table 5. Abbreviations

Name	Object type
Host	Refers to a single host
Host cluster	Refers to a host cluster (which is part of a servers that share a set of volumes)
Host object	Refers to a list of worldwide port names WWPNs and iSCSI names
Volume copy	Referred to as vdiskcopy
Thin-provisioned volume copy	Referred to as sevdiskcopy
Managed disk	Referred to as MDisk
MDisk group or storage pool	Referred to as storage pool
I/O group	Referred to as iogrp
Node, node canister, enclosure	Refers to node or node canister or enclosure, depending on the system
Clustered system (system) or cluster	Referred to as system
VDisk or volume	Referred to as volume
Controller	Refers to a controller that is part of a clustered system
IBM FlashCopy® mapping	Referred to as fcmap
IBM FlashCopy consistency group	Referred to as fcconsistgrp
Metro Mirror or Global Mirror relationship	Referred to as rrelationship

Table 5. Abbreviations (continued)

Name	Object type
Metro Mirror or Global Mirror consistency group	Referred to as rconsistgrp
HyperSwap [®] mirroring consistency group	Referred to as hyperswap
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.

delimiters (: or , or !)

Delimiters are used to delimit items listed after issuing an information command.

- Colon (:) is used to delimit items in a list in a command (for example `mkhost -name myhost -hbawpwn AA22000011112222:AA22000011112223`).
- Comma (,) is used to delimit items in a list in a command if item values can contain a colon.
- Exclamation mark (!) is used to delimit items in a command if the item values can contain a colon or comma. Exclamation points generally do not show up in example output and are a good delimiter to use.

Using wildcards in the CLI

You can use wildcards in the SAN Volume Controller command-line interface (CLI).

The CLI 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:

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

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

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 clustered system (system). 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 system shell, this will result in the wildcard being expanded in the system shell.

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

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

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

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

Data types and value ranges

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

Note: When creating a new object, the clustered system (system) assigns a default **-type** name if one is not specified. The default **-type** name consists of the object prefix and the lowest available integer starting from 0 (except for nodes starting from 1); for example, *vdisk23*; the default -type name must be unique.

Table 6 lists the data types and the value ranges for each.

Table 6. Data types

Data types	Value ranges
filename_arg	This is a (optionally fully qualified) file name, containing a maximum of 169 characters. Valid characters are: <ul style="list-style-type: none">• . (period; the field must not start with, end with, or contain two consecutive periods)• / (forward slash)• - (hyphen)• _ (underscore)• a–z (lowercase letters, A through Z)• A–Z (uppercase letters, A through Z)• 0–9 (numerals 0 through 9)

Table 6. Data types (continued)

Data types	Value ranges
directory_or_file_filter	<p>Specifies a directory, file name filter, or both, within the specified directory. Valid directory values are:</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, containing a maximum of 128 characters, with or without the "*" (wildcard), and appended to the end of a directory value. Valid characters are:</p> <ul style="list-style-type: none"> • * (asterisk/wildcard) • . (the field must not start with, end with, or contain two consecutive periods) • / • - • _ • a-z • A-Z • 0-9
filename_prefix	<p>The prefix of a file name, containing a maximum of 128 characters. Valid characters are:</p> <ul style="list-style-type: none"> • - • _ • a-z • A-Z • 0-9
name_arg	<p>Names can be specified or changed using the create and modify functions. The view commands provide both the name and ID of an object.</p> <p>Note: The system name is set when the system is created.</p> <p>The first character of a <i>name_arg</i> must be nonnumeric. The first character of an object name cannot be a - (dash) because the CLI (command-line interface) interprets it as being the next parameter.</p> <p>Valid characters are:</p> <ul style="list-style-type: none"> • . (a period - the field must not start with, end with, or contain two consecutive periods) • / • - • _ • space • a through z • A through Z • 0 through 9

Table 6. Data types (continued)

Data types	Value ranges
password	<p>This is a user-defined password containing a maximum of 15 characters. Valid characters are:</p> <ul style="list-style-type: none"> • - (cannot be used as the first character) • _ • a–z • A–Z • 0–9
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.</p> <p>The standard defines a way to extend the serial number using letters in the place of numbers in the 5-digit field.</p>
ip_address_arg	<p>The argument follows the standard rules for dotted decimal notation.</p> <p>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	<p>This is the dotted domain name for the system subnet (for example, ibm.com).</p>
hostname	<p>The host name assigned to the system. This name can be different from the system name, and is modifiable.</p> <p>A combination of the host name and the <i>dns_name</i> is used to access the system, for example: https://hostname.ibm.com/</p>
capacity_value	<p>The capacity expressed within a range of 512 bytes to 2 petabytes (PB).</p> <p>Tip: Specify the capacity as megabytes (MB), kilobytes (KB), gigabytes (GB), or PB. When using MB, specify the value in multiples of 512 bytes. A capacity of 0 is valid for a striped or sequential volume. The smallest number of supported bytes is 512.</p>
node_id	<p>A node ID differs from other IDs in that it is a unique ID assigned when a node is used to create a system, or when a node is added to a system. A <i>node_id</i> value is never reused in a system.</p> <p>Node IDs are internally represented as 64-bit numbers, and like other IDs, cannot be modified by user commands.</p>

Table 6. Data types (continued)

Data types	Value ranges
xxx_id	<p>All objects are referred to by unique integer IDs, 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 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 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), expressed as a 64-bit hexadecimal number and consisting of the characters 0–9, a–f, and A–F; for example: 1A2B30C67AFFE47B.</p> <p>Note: Entering WWPN 0 in the command string causes a command failure.</p>
panel_name	This is a string of up to six characters corresponding to the number on the printed label below the display on the front panel of a node in the system.
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 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.

Table 6. Data types (continued)

Data types	Value ranges
directory_arg	<p>Specifies a directory, file name filter, or both, within the specified directory. Valid directory values are:</p> <ul style="list-style-type: none"> • /dumps • /dumps/audit • /dumps/cimom • /dumps/configs • /dumps/elog • /dumps/feature • /dumps/iostats • /dumps/iotrace • /home/admin/upgrade <p>The file name filter can be any valid file name, containing a maximum of 128 characters, with or without the wildcard (*, an asterisk), and appended to the end of a directory value. Valid characters are:</p> <ul style="list-style-type: none"> • * • . (the field must not start with, end with, or contain two consecutive periods) • / • - • _ • a-z • A-Z • 0-9
locale_arg	<p>The system locale setting. Valid values are:</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 a secure shell (SSH) key, containing a maximum of 30 characters.
user_arg	Specifies the user: 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 entered into a colon-separated list is 128; exceeding this maximum number returns an error.

CLI commands and parameters

Command-line interface (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:

- When a command name is specified, the first argument given must be the action that you want to be performed.
- Where you are specifying a command 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.
chuser -usergrp=-usergrp -- -password
- The symbol '--' is valid as the final word on the command line.

Examples that are valid

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

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

Examples that are invalid

```
chuser -usergrp 1 fred --
chuser -usergrp 1 -- fred --
chuser -- -usergrp 1 fred
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 **lssystem -h** provides a list of the actions available with the **lssystem** command.

-nomsg

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

```
mkmdiskgrp -ext 16
```

it displays:

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

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

```
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.

CLI deprecated and discontinued commands

Some command-line interface (CLI) commands are discontinued or deprecated and replaced with new commands.

Command-line interface (CLI) commands can be discontinued or deprecated and replaced with new commands that are more effective. This action can affect scripting, which is done to simplify tasks such as specifying (repeated) commands.

A discontinued command is removed from the CLI and can no longer be used. In most instances, there is a replacement command for discontinued commands.

A deprecated command can still be used if the appropriate command prefix is specified (`svctask` or `svcinfo`, for example). In most instances, there is a replacement command for deprecated commands.

The following commands are deprecated:

- These commands are replaced by **lseventlog**:

- **caterrlog**
- **caterrlogbyseqnum**
- **lserrlogbyfcconsistgrp**
- **lserrlogbyfcmap**
- **lserrlogbyhost**
- **lserrlogbyiogrp**
- **lserrlogbymdisk**
- **lserrlogbymdiskgrp**
- **lserrlogbynode**
- **lserrlogbyrcconsistgrp**
- **lserrlogbyrcrelationship**
- **lserrlogbyvdisk**

- These commands are replaced by **lsdumps**:

- **ls2145dumps**
- **lsauditlogdumps**

- **lserrlogdumps**
- **lscimomdumps**
- **lsfeaturedumps**
- **lsiostatsdumps**
- **lsiotracedumps**
- **lsdiskdumps**
- **lssoftwaredumps**
- **lssoftwareupgradestatus** is replaced by **lsupdate**
- **chenclosurevpd** is replaced by **chvpd**
- **cherrstate** is replaced by **cheventlog**
- **lsnodependentvdisks** is replaced by **lsdependentvdisks**
- **setquorum** is replaced by **chquorum**
- **mkpartnership** is replaced by **mkippartnership** and **mkfcpartnership**
- **lshbaportcandidate** is replaced by **lssasportcandidate** and **lsfcportcandidate**

The following commands are discontinued:

- These commands are replaced by **lsdumps**:
 - **svcservicemodeinfo ls2145dumps**
 - **svcservicemodeinfo lsclustervpd**
 - **svcservicemodeinfo lserrlogdumps**
 - **svcservicemodeinfo lsfeaturedumps**
 - **svcservicemodeinfo lsiostatsdumps**
 - **svcservicemodeinfo lsiotracedumps**
 - **svcservicemodeinfo lsdiskdumps**
 - **svcservicemodeinfo lssoftwaredumps**
- These commands are replaced by the **user management commands**:
 - **addsshkeys**
 - **lsauth**
 - **mkauth**
 - **rmsshkey**
 - **rmlsshkeys**
 - **rmauth**
- **applydisksoftware** is replaced by **applydrivesoftware**
- **chcluster** is replaced by **chsystem**
- **cpfabricdumps** has no replacement
- **dumpconfig** has no replacement
- **dumpinternallog** has no replacement
- **lscluster** is replaced by **lsystem**
- **lsclustercandidate** is replaced by **lspartnershipcandidate**
- **lsclusterip** is replaced by **lssystem**
- **lsclusterstats** is replaced by **lssystemstats**
- **lsconfigdumps** has no replacement
- **recoverarraybycluster** is replaced by **recoverarraybysystem**
- **recoverdiskbycluster** is replaced by **recoverdiskbysystem**
- **svcservicemodeinfo lsclustervpd** is replaced by **satask lsservicestatus**

- **svcservicemodetask applysoftware** is replaced by **satask installsoftware**
- **svcservicemodetask cleardumps** is replaced by **cleardumps**
- **svcservicemodetask dumperrlog** is replaced by **dumperrlog**
- **svcservicemodetask exit** is replaced by **stopservice**
- **setclustertime** is replaced by **setsystemtime**
- **stopcluster** is replaced by **stopsystem**
- **triggermdiskdump** is replaced by **triggerdrivedump**
- **setevent** is replaced by **chsnmpserver**, **lssnmpserver**, **mksnmpserver**, and **rmsnmpserver**
- **setemail** is replaced by **chemail**, **chemailserver**, **lsemailserver**, **mkemailserver**, and **rmemailserver**

Understanding capacity indicators

The system uses base-2 (binary numeral) as capacity indicators for volumes, drives, and other system objects. The management GUI and the command-line interface (CLI) use different abbreviations to indicate capacity, but the value for these capacity indicators is the same.

The following table displays the differences in how capacity indicators are displayed in the management GUI and the CLI.

Table 7. Capacity indicators. This table displays the differences in how capacity indicators are displayed in the management GUI and the CLI.

Metric	GUI Abbreviation	CLI Abbreviation	Value
kibibyte	KiB	KB	1024
mebibyte	MiB	MB	1,048,576
gibibyte	GiB	GB	1,073,741,824
tebibyte	TiB	TB	1,099,511,627,776
pebibyte	PiB	PB	1,125,899,906,842,624
exbibyte	EiB	EB	1,152,921,504,606,846,976
zebibyte	ZiB	ZB	1,180,591,620,717,411,303,424
yobibyte	YiB	YB	1,208,925,819,614,629,174,706,176

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 less than (<) and greater than (>) 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)
- kb (Kilobytes)
- 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.

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.

Chapter 1. Setting up an SSH client

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

Overview

The system acts as the SSH server in this relationship. The SSH client provides a secure environment in which to connect to a remote machine. Authentication is completed using user name and password. If you require command-line access without entering a password, it uses the principles of public and private keys for authentication.

Authenticating SSH logins

Generate a Secure Shell (SSH) key pair to use the command-line interface (CLI). Additionally, when you use the SSH to log in to the system, you must use the RSA-based private key authentication.

When you are using AIX® hosts, SSH logins are authenticated on the system using the RSA-based authentication that is supported in the OpenSSH client that is available for AIX. This scheme is based on the supplied password (or if you require command-line access without entering a password, then public-key cryptography is used) by using an algorithm known commonly as RSA.

Note: The authentication process for host systems that are not AIX 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 system, 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 SSH password (if you require command-line access without entering a password, 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).

The system supports up to 32 interactive SSH sessions on the management IP address simultaneously.

Note: After one hour, a fixed SSH interactive session times out, which means the SSH session is automatically closed. This session timeout limit is not configurable.

To connect to the system, the SSH client requires a user login name and an SSH password (or if you require command-line access without entering a password, the key pair). Authenticate to the system using a management user name and password. When using an SSH client to access a system, you must use your *SVC_username* and password. The system uses the password (and if not a password, the SSH key pair) to authorize the user accessing the system.

You can connect to the system using the same user name with which you log into the system.

For Microsoft Windows hosts, PuTTY can be downloaded from the Internet and used at no charge to provide an SSH client.

You can connect to the system using the same user name with which you log into the system.

Setting up an SSH client on a Windows host

You can prepare the SSH client on a Windows host.

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 workstation in one of these ways:

- If you purchased 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.
- You can use the separate PuTTY client program-installation wizard, *putty-version-installer.exe*. You can download the PuTTY client program from this website:

Download Putty

Note: Before you install the PuTTY client program, ensure that your Windows system meets the system requirements.

To connect to the system, the SSH client requires a user login name and an SSH password (or if you require command-line access without entering a password, the key pair). Authenticate to the system using a management user name and password. When using an SSH client to access a system, you must use your *SVC_username* and password. The system uses the password (and if not a password, the SSH key pair) to authorize the user accessing the system.

You can connect to the system using the same user name with which you log into the system.

Generating an SSH key pair using PuTTY

To use the SAN Volume Controller command-line interface, you must generate a Secure Shell (SSH) key pair using PuTTY.

About this task

Generate SSH keys using the PuTTY key generator (PuTTYgen):

Procedure

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:
 - 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:
 - 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) password. If you require command line access without entering a password, use the SSH key pair that you created for the command-line interface (CLI).

About this task

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.

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

Procedure

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 password. If no password is used, the private key in the **Private key file for authentication** field.
9. Click **Connection > Data** in the Category navigation tree.
10. Type the user name that you want to use on the SAN Volume Controller 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 Internet Protocol (IP) address of one of the SAN Volume Controller clustered system (system) IP addresses or host names.
13. Type **22** in the **Port** field. The SAN Volume Controller system 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 System 1.

15. Click **Save**.

Results

You have now configured a PuTTY session for the CLI.

Note: If you configured more than one IP address for the SAN Volume Controller system, 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 website: [Download Putty](#).

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 or user password. 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 clustered system. 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 system supports up to 32 interactive SSH sessions on the management IP address simultaneously.

Note: After one hour, a fixed SSH interactive session times out, which means the SSH session is automatically closed. This session timeout limit is not configurable.

The following example shows the command that you can issue to start interactive mode:

```
C:\support utils\putty <username>@svcconsoleip
```

where *support utils\putty* is the location of your putty.exe file, *<username>* is the IP address of your management GUI, and *<username>* is the user name that you want to use.

If you were to issue the **lsuser** command, which lists the SSH client public keys that are stored on the system, the following output is displayed when *ssh_key=yes*:

```
IBM_2145:cluster0:superuser>lsuser
id name      password ssh_key remote usergrp_id usergrp_name
0  superuser yes      yes      no      0      SecurityAdmin
1  smith     no       yes      no      4      Monitor
2  jones     no       yes      no      2      CopyOperator
```

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\svconconsole\cimom>plink superuser@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 user name "superuser".
Authenticating with public key "imported-openssh-key"
IBM_2145:your_cluster_name:superuser>
```

Single line command

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

```
C:\Program Files\IBM\svconconsole\cimom>
plink superuser@9.43.225.208 lsuser
Authenticating with public key "imported-openssh-key"
id name      password ssh_key remote usergrp_id usergrp_name
0  superuser yes      yes      no      0      SecurityAdmin
1  smith     no       yes      no      4      Monitor
2  jones     no       yes      no      2      CopyOperator
```

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\svconsole\cimom>
plink superuser@9.43.225.208 lsuser
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"
id name      password ssh_key remote usergrp_id usergrp_name
0  superuser  yes      yes      no       0          SecurityAdmin
1  smith      no       yes      no       4          Monitor
2  jones      no       yes      no       2          CopyOperator

```

Starting a PuTTY session for the CLI

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

Before you begin

This task assumes that you have already configured and saved a PuTTY session using the Secure Shell (SSH) password. If you require command line access without entering a password, use the SSH key pair that you created for the CLI.

About this task

Start a PuTTY session:

Procedure

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 password or key pair, a PuTTY Security Alert window is displayed. Click **Yes** to accept the change and trust the new key.

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

Preparing the SSH client on an AIX or Linux host

You can prepare the Secure Shell (SSH) client on an AIX or Linux host.

Before you begin

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 POWER 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 website:

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 website:

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 website:

www.openssh.org/portable.html

About this task

Authentication to the system generally requires the use of a password, but if there is no password you can use a key pair. Use these steps to set up an RSA key pair on the AIX or Linux host and the clustered system:

Results

To authenticate using an SSH key, use the following command:

```
ssh -i full_path_to_key username@my_system
```

Where *my_system* is the name of the system IP, *username@my_system* is the user name that you also log into the system with, and *full_path_to_key* is the full path to the *key* file that was generated in the previous step. Authenticate to the system using a *SVC_username* and password. (If you require command-line access without using a password, SSH keys can be used.) The system determines which user is logging in from the key the user is using.

Note: You can omit `-i full_path_to_key` if you configure the SSH client to use the key file automatically.

If you use the Secure Shell (SSH) to log into the system, use the password defined for accessing the GUI. You can also use RSA-based private key authentication.

For more information, see “Connecting to the CLI using OpenSSH” on page 8.

Generating an SSH key pair using OpenSSH

You can generate an SSH key pair using OpenSSH.

About this task

Set up an RSA key pair on the AIX or Linux host and the clustered system:

Procedure

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

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

You can also create a valid ECDSA key pair for authentication:

```
ssh-keygen -t ecdsa
```

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 clustered system using the management GUI.

Connecting to the CLI using OpenSSH

You can connect to the command-line interface (CLI) using OpenSSH.

To connect to a clustered system using a *SVC_username* and SSH password, issue:

```
ssh username@my_system
```

To use an SSH key, issue:

```
-i full_path_to_key
```

Where *my_system* is the name of the system IP, *full_path_to_key* is the full path to the key file that was generated in the previous step, and *SVC_username* is the user name that you want to use on SAN Volume Controller.

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

Working with local and remote users

You can create either a local or a remote user to access a SAN Volume Controller clustered system (system).

Before you begin

You can create two categories of users that access the system. These types are based on how the users are authenticated to the system. Local users must provide the *SVC_username* and password, and if you require command line access without entering 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 system.

If the local user needs access to management GUI, a password is needed for the user. Access to the command-line interface (CLI) is also possible with the same password or (alternatively) a valid SSH key can be used. An SSH password is required if a user is working with both interfaces. User groups define roles that authorize the users within that group to a specific set of operations on the system.

Local users must be part of a user group that is defined on the system.

A remote user is authenticated on a remote service usually provided by a SAN management application, such as IBM Spectrum Control, and does not need local authentication methods. For a remote user, a password (preferred) is required, and if you require command line access without entering a password an SSH key is required to use the command-line interface.

Remote users only need local credentials to access the management GUI if the remote service is down. The user groups a remote user is a member of are defined by the remote authentication service. To define a remote user, create an user group on the local machine that is also defined on the remote authentication service.

You can connect to the system using the same user name with which you log into the system.

About this task

Complete these steps to create either a local or remote user:

Procedure

1. Select **Access > Users** .
2. Select the appropriate user group.

3. Click **Create User** .
4. Enter the information on the new user and click **Create**.

UNIX commands available in interactive SSH sessions

You can use several UNIX-based commands while working with interactive SSH sessions.

The system supports up to 32 interactive SSH sessions on the management IP address simultaneously.

Note: After one hour, a fixed SSH interactive session times out, which means the SSH session is automatically closed. This session timeout limit is not configurable.

You can use the following UNIX commands to manage interactive SSH sessions:

Table 8. UNIX commands for interactive SSH sessions

UNIX commands	Description
grep	Filters output by keywords or expressions.
more	Moves through output one page at a time.
sed	Filters output by complex expressions.
sort	Sorts output according to criteria.
cut	Removes individual columns from output.
head	Displays only first lines.
less	Moves through the output bidirectionally a page at a time. (secure mode)
tail	Shows only the last lines.
uniq	Hides any duplicate information.
tr	Translate characters.
wc	Counts lines and words and characters in data.

Copying the software update files by using PuTTY pscp or openssh scp

PuTTY pscp (or scp) 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.

Before you begin

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

About this task

The pscp or scp 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. For Linux users, the scp is installed with the openssh packages.

Complete these steps to use the pscp application. The scp process is similar from step 6 on page 10.

Procedure

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

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

where C:\Program Files\putty is the directory where PuTTY is installed.

6. Issue this 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  
username@cluster_ip_address:/home/admin/update
```

where *saved_putty_configuration* is the name of the PuTTY configuration session, *directory_software_upgrade_files* is the location of the software update files, *software_upgrade_file_name* is the name of the software update file, *username* is the name that you want to use on the SAN Volume Controller, and *cluster_ip_address* is an IP address of your clustered system.

Note: Saving the PuTTY configuration session in step 3 and then loading the PuTTY configuration session in step 6 is optional. To copy without loading a PuTTY configuration session, use the following syntax:

```
pscp directory_software_upgrade_files/software_upgrade_file_name  
username@cluster_ip_address:/home/admin/update
```

If there is insufficient space to store the software update file on the system, the copy process fails. In this case, complete the following steps:

- a. Use `pscp` to copy data that you want to preserve from the `/home/admin/update` directory.
- b. Use the following command to delete dump files in the `/home/admin/update` directory:

```
cleardumps -prefix /home/admin/update
```
- c. Repeat step 6.

Chapter 2. Using the CLI

The SAN Volume Controller clustered system (system) 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 system.

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

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.
- Authenticate to the system using a password.
- If you require command line access without entering a password, use an SSH public key. Then 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 for these functions:

- Setup of the system, its nodes, and the I/O groups
- Analyze error logs (event logs)
- Setup and maintenance of managed disks (MDisk) and storage pools
- Setup and maintenance of client public SSH keys on the system
- Setup and maintenance of volumes
- Setup of logical host objects
- Map volumes to hosts
- Navigate from managed hosts to volumes 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 clustered system time by using the CLI

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

About this task

To set the system time:

Procedure

1. Issue the **showtimezone** CLI command to display the current time-zone settings for the system. The time zone and the associated time-zone ID are displayed.

2. Issue the **ls timezones** CLI command to list the time zones that are available on the system. 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 system.

```
settimezone -timezone time_zone_setting
```

where 031809142005*time_zone_setting* is the new time zone ID that you chose from the list of time zones that are available on the system.
4. Issue the following CLI command to set the time for the system:

```
setsystemtime -time 031809142005
```

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

Setting cluster date and time

You can set the date and time for a SAN Volume Controller cluster from the System Date and Time Settings panel.

Before you begin

This task assumes that you already launched the management GUI.

About this task

You can set the System Date and time manually, or by specifying an NTP server:

Procedure

1. Click **Manage Systems > Set System Time** in the portfolio. The System Date and Time Settings panel is displayed.
2. To use NTP to manage the clustered system date and time, enter an Internet Protocol Version 4 (IPv4) address and click **Set NTP Server**.

Note: If you are using a remote authentication service to authenticate users to the SAN Volume Controller clustered system, then both the system and the remote service should use the same NTP server. Consistent time settings between the two systems ensure interactive performance of the management GUI and correct assignments for user roles.

3. To set the clustered system date and time manually, continue with the following steps.
4. Type your changes into the **Date**, **Month**, **Year**, **Hours**, and **Minutes** fields and select a new time zone from the **Time Zone** list.
5. Select **Update cluster time and date**, **Update cluster time zone**, or both.
6. Click **Update** to submit the update request to the clustered system.

Viewing and updating license settings by using the CLI

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

About this task

SAN Volume Controller provides two license options: Physical Disk Licensing and Capacity Licensing. To view and update your SAN Volume Controller license settings:

Procedure

1. Issue the **ls license** CLI command to view the current license settings for the clustered system (system).

2. Issue the **chlicense** CLI command to change the licensed settings of the system.

Attention:

- License settings are entered when the system is first created; do not update the settings unless you change your license.
- To select Physical Disk Licensing, run the **chlicense** command with one or more of the **physical_disks**, **physical_flash**, and **physical_remote** parameters.
- To select Capacity Licensing, run the **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 clustered system properties by using the CLI

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

About this task

These actions help you display your system property information.

Procedure

Issue the **lssystem** command to display the properties for a system.

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

```
lssystem -delim : build1
```

where *build1* is the name of the system.

Results

```
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_alterate:
email_contact_location:
email_state:stopped
inventory_mail_interval:0
total_vdiskcopy_capacity:15.71GB
total_used_capacity:13.78GB
total_overalllocation: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_alterate:
total_allocated_extent_capacity:16.12GB
```

Maintaining passwords using the CLI

You can use the command-line interface (CLI) to view and change the status of the password reset feature for the system.

The superuser password can be reset to its default value of **passw0rd** by using the technician port on SAN Volume Controller 2145-DH8 nodes or the front panel on earlier models of the system. 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 **setpwdreset** CLI command to view and change the status of the password reset feature for the system.
2. Record the system superuser password because you cannot access the system without it.

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

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

Using the dump commands to work with directories

The **lsdumps** 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
- /dumps/mdisk
- /home/admin/update
- /dumps/drive
- /dumps/enclosure

Use the **lsdumps** 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. If you do not specify a node, the available dumps on the configuration node are listed.

Use the **cpdumps** command to copy dump files from a nonconfiguration node to a the configuration node. You can use this command to retrieve dumps that were saved to an older configuration node. You can retrieve files and put them on the configuration node to be copied.

Use the **cleardumps** command to delete dump directories on a specified node. You can clear specific files or groups of files based on the use of a wildcard (an asterisk, *). You can delete files on a single directory or all of the dump directories (by specifying the /dumps variable).

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 **catauditlog** command. To dump the contents of the audit log to a file on the current configuration node, issue the **dumpauiditlog** 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 clustered system (system). 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 **dumpperrlog**

command. This dumps the contents of the error or event log to the /dumps/elog 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 **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 **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 system. An I/O statistics dump is created by using the **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, volume, 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 **settrace** command. The collection of the I/O trace data is started by using the **starttrace** command. The I/O trace data collection is stopped when the **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 **settrace** command, and `NNNNNN` is the node name.

Dumps that are contained in the /dumps/mdisk directory are copies of flash drive MDisk internal logs. These dumps are created using the **triggerdrivedump** command. The file name is `mdiskdump_NNNNNN_MMMM_YYMMDD_HHMMSS`, where `NNNNNN` is the name of the node that contains the MDisk, and `MMMM` is the decimal ID of the MDisk.

Software update packages are contained in the /home/admin/upgrade directory. These directories exist on every node in the system.

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 might 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 **cpdumps** command to copy the files from a nonconfiguration node to the current configuration node.

Re-adding a repaired node to a clustered system by using the CLI

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

Before you begin

Before you add a node to a clustered system, you must make sure that the switchd\ zoning is configured such that the node that is being added is in the same zone as all other nodes in the clustered system. 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 that is 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 select the correct I/O group can result in data corruption. You must use the information that was recorded when the node was originally added to the clustered system. If you do not have access to this information, call the IBM Support Center to add the node back into the clustered system 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 clustered system. You must ensure that the LUNs are the same before you add the new node to the clustered system.
3. LUN masking for each LUN must be identical on all nodes in a clustered system. You must ensure that the LUN masking for each LUN is identical before you add the new node to the clustered system.
4. You must ensure that the model type of the new node is supported by the SAN Volume Controller software level that is installed on the clustered system. If the model type is not supported by the SAN Volume Controller software level, update the clustered system to a software level that supports the model type of the new node. See the following website for the latest supported software levels:
www.ibm.com/support

About this task

Special procedures when you add a node to a clustered system

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 that are supported by the Subsystem Device Driver (SDD). SDD maintains an association between a vpath and a SAN Volume Controller volume. This association uses an identifier (UID) which is unique to the volume and is never reused. The UID permits SDD to directly associate vpaths with volumes.

SDD operates within a protocol stack that contains disk and Fibre Channel device drivers that are used to communicate with the SAN Volume Controller using the SCSI protocol over Fibre Channel as defined by the ANSI FCS standard. The addressing scheme that is 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 by using the same WWNN and LUN numbers that were previously used.

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

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

- The clustered system has more than one I/O group.
- The node that is being added to the clustered system uses physical node hardware or a slot that has previously been used for a node in the clustered system.

- The node that is being added to the clustered system uses physical node hardware or a slot that has previously been used for a node in another clustered system and both clustered systems 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 **lsnode** or the management GUI to determine the WWN of the clustered system nodes.
- Before you add the node back into the clustered system, you must shut down all of the hosts using the clustered system. 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 by using the clustered system, then do the following:
 - On all of the hosts that are connected to the clustered system, unconfigure the Fibre Channel adapter device driver, the disk device driver, and multipathing driver before you add the node to the clustered system.
 - Add the node to the clustered system, 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 clustered system 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 clustered system by using the CLI command **addnode** or the management GUI.

Note: You do not need to run the **addnode** command on a node with a partner that is already in a clustered system; the clustered system automatically detects an online candidate.

Note: The **addnode** command is a SAN Volume Controller command. For Storwize V7000, use the **addcontrolenclosure** command.

- A user decides to delete four nodes from the clustered system and add them back into the clustered system using the CLI command **addnode** or the management GUI.

Note: The **addnode** command is a SAN Volume Controller command. For Storwize V7000, use the **addcontrolenclosure** command.

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

Procedure

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

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

```
lsnode -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 example:

```
lsnode -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
```

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

Note: The **lsnodecandidate** command is a SAN Volume Controller command. For Storwize V7000, use the **lscontrolenclosurecandidate** command.

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

```
lsnodecandidate -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 **addnode** CLI command to add a node to the clustered system.

Note: The **addnode** command is a SAN Volume Controller command. For Storwize V7000, use the **addcontrolenclosure** command.

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 clustered system by using the panel name parameter:

```
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 clustered system by using the WWNN parameter:

```
addnode -wwnname 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 **lsnode** CLI command to verify the final configuration.

The following example shows output that is displayed:

```
lsnode -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.ndih11.node1:
```

Record the following information for the new node:

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

Note: If this command is issued quickly after you add nodes to the clustered system, the status of the nodes might be adding. The status is shown as adding if the process of adding the nodes to the clustered system 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.

Results

The nodes are added to the clustered system.

Displaying node properties by using the CLI

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

About this task

To display the node properties:

Procedure

1. Use the **lsnode** CLI command to display a concise list of nodes in the clustered system.

Issue this CLI command to list the system nodes:

```
lsnode -delim :
```

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

The following example is a CLI command that you can use to list detailed output for a node in the system:

```
lsnode -delim : group1node1
```

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

Discovering MDisks using the CLI

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

About this task

The clustered system (system) automatically discovers the back-end controller and integrates the controller to determine the storage that is presented to the SAN Volume Controller nodes when back-end controllers are:

- Added to the Fibre Channel
- Included in the same switch zone as a SAN Volume Controller system

The Small Computer System Interface (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 system might be unaware of these configuration changes. You can request that the SAN Volume Controller system rescans the Fibre Channel SAN to update the list of unmanaged MDisks.

Note: The automatic discovery completed by the system does not write anything to an unmanaged MDisk. You must instruct the system to add an MDisk to a storage pool or use an MDisk to create an image mode volume.

Discover (and then view) a list of MDisks:

Procedure

1. Issue the **detectmdisk** CLI command to manually scan the Fibre Channel network. The scan discovers any new MDisks that might have been added to the system and can help rebalance MDisk access across the available controller device ports.

Notes:

- a. Only issue the **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 **lsmdiskcandidate** CLI command to show the unmanaged MDisks. These MDisks have not been assigned to a storage pool.
 3. Issue the **lsmdisk** CLI command to view all of the MDisks.

Results

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

Example

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

1. Issue **detectmdisk**.
2. Issue **lsmdiskcandidate**.

This output is displayed:

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

3. Issue **lsmdisk -delim :**

This output is displayed:

```
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:
20000004cf2422aa0000000000000000000000000000000000000000000000000000:
1:mdisk1:online:unmanaged:::68.4GB:0000000000000000:controller1:
20000004cf1fd19d0000000000000000000000000000000000000000000000000000:
2:mdisk2:online:unmanaged:::68.4GB:0000000000000000:controller2:
20000004cf242531000000000000000000000000000000000000000000000000000:
```

Creating storage pools using the CLI

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

Before you begin

Attention: If you add an MDisk to a 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 volumes instead.

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

If you are using a SAN Volume Controller flash drive managed disk, ensure that you are familiar with the flash drive configuration rules.

If you intend to keep the volume allocation within one storage system, ensure that all MDiskS in the storage pool are presented by the same storage system.

Ensure that all MDisks that are allocated to a single storage pool are of the same RAID type. If the 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 a storage pool in the same tier must be similar and have similar performance characteristics. If you do not use Easy Tier, the storage pool must contain only one tier of storage, and all of the MDisks in the storage pool must be similar and have similar performance characteristics.

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

- A volume can only be created using the storage from one storage pool. Therefore, if you create small (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 storage pool goes offline, all the (volumes) in the storage pool go offline. Therefore you might want to consider using different 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 storage pool.
- All the MDiskS in a storage pool must have similar levels of performance or reliability, or both. If a storage pool contains MDiskS with different levels of performance, the performance of the (volumes) in this group is limited by the performance of the slowest MDisk. If a storage pool contains MDiskS with different levels of reliability, the reliability of the (volumes) in this group is that of the least reliable MDisk in the group.

Note: When you create a pool with a new flash drive, the new flash drive is automatically formatted and set to a block size of 512 bytes.

About this task

Even with the best planning, circumstances can change and you must reconfigure your (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 a storage pool extent size

Consider the following factors as you plan the extent size of each new storage pool:

- You must specify the extent size when you create a new storage pool.
- You cannot change the extent size later; it must remain constant throughout the lifetime of the storage pool.
- Storage pools can have different extent sizes; however, this places restrictions on the use of data migration.
- The extent size affects the maximum size of a volume in the storage pool. A larger extent size increases the total amount of storage that the system can manage, and a smaller extent size allows more fine-grained control of storage allocation.

Table 9 compares the maximum 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 volume that is created, using a larger extent size might increase the amount of storage that is wasted at the end of each 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 9. Maximum volume capacity by extent size

Extent size (MB)	Maximum volume capacity in GB (not thin-provisioned volumes)	Maximum 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 storage pools; however, you cannot migrate (volumes) between storage pools with different extent sizes. If possible, create all your storage pools with the same extent size.

Use the following steps to create a storage pool:

Procedure

Issue the **mkmdiskgrp** CLI command to create a storage pool.

This is an example of the CLI command you can issue to create a storage pool:

```
mkmdiskgrp -name maindiskgroup -ext 32  
-mdisk mdsk0:mdsk1:mdsk2:mdsk3
```

where *maindiskgroup* is the name of the 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 MDisk that you want to add to the group.

Results

You created and added MDisk to a storage pool.

Example

The following example provides a scenario where you want to create a storage pool, but you do not have any MDisk available to add to the group. You plan to add the MDisk at a later time. You use the **mkmdiskgrp** CLI command to create the storage pool *bkpmdiskgroup* and later used the **addmdisk** CLI command to add *mdsk4*, *mdsk5*, *mdsk6*, *mdsk7* to the storage pool.

1. Issue **mkmdiskgrp -name bkpmdiskgroup -ext 32**
where *bkpmdiskgroup* is the name of the 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 MDisk that you want to add to the storage pool.
3. Issue **addmdisk -mdisk msk4:msk5:msk6:msk7 bkpdiskgroup**
where *mdsk4*, *mdsk5*, *mdsk6*, *mdsk7* are the names of the MDisk that you want to add to the storage pool and *bkpdiskgroup* is the name of the storage pool for which you want to add MDisk.

Adding MDisk to storage pools using the CLI

You can use the command-line interface (CLI) to add managed disk (MDisk) to storage pools.

Before you begin

The MDisk must be in unmanaged mode. Disk that already belong to a storage pool cannot be added to another storage pool until they have been deleted from their current storage pool. An MDisk can be deleted from a storage pool under these circumstances:

- If the MDisk does not contain any extents in use by a volume
- If you can first migrate the extents in use onto other free extents within the group

About this task

Important: Do not add an MDisk using this procedure if you are mapping the MDisk to an image mode volume. Adding an MDisk to a storage pool enables 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 volume, use the **mkvdisk** command instead of **addmdisk**.

If you are using a SAN Volume Controller flash drive managed disk, ensure that you are familiar with the flash drive configuration rules.

The SAN Volume Controller performs tests on the MDisk in the list before the MDisk are allowed to become part of a storage pool when:

- Adding MDisk to a storage pool using the **addmdisk** command
- Creating a storage pool using the **mkmdiskgrp -mdisk** command

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 MDisk are not added to the group. However, with the **mkmdiskgrp -mdisk** command, the storage pool is still created even if the tests fail, but it does not contain any MDisk. If tests fail, confirm that the MDisk are in the correct state and that they have been correctly discovered.

These events contribute to an MDisk test failure:

- The MDisk is not visible to all SAN Volume Controller nodes in the clustered system.
- 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.

These 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 flash drive to a pool, the flash drive is automatically formatted and set to a block size of 512 bytes.

Procedure

To add MDisk to storage pools, complete the following steps.

1. Issue the **lsmdiskgrp** CLI command to list the existing storage pools.

This example is a CLI command that you can issue to list the existing storage pools:

```
lsmdiskgrp -delim :
```

This 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 **addmdisk** CLI command to add MDisk to the storage pool.

This is an example of the CLI command you can issue to add MDisk to a storage pool:

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

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

Setting a quorum disk using the CLI

You can set an external managed disk (MDisk) as a quorum disk by using the command-line interface (CLI).

Note: Quorum functionality is not supported for internal drives on SAN Volume Controller nodes.

To set an MDisk as a quorum disk, use the **chquorum** command. **Storwize V7000:** To set an external MDisk as a quorum disk, use the **chquorum** command.

When setting an MDisk as a quorum disk, keep the following recommendations in mind:

- 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, search for supported hardware list at the following website:

www.ibm.com/support

- Before you set the quorum disk with the **chquorum** command, use the **lsquorum** command to ensure that the MDisk you want is online.

Quorum disk configuration describes how quorum disks are used by the system, and how they are selected. The system automatically assigns quorum disks. Do not override the quorum disk assignment if you have a system without external MDisks. For a system with more than one control enclosure and with external MDisks, distribute the quorum candidate disks (when possible) so that each MDisk is provided by a different storage system. For a list of storage systems that support quorum disks, search for supported hardware list at the following website:

www.ibm.com/support

Modifying the amount of available memory for Copy Services, Volume Mirroring, and RAID arrays by using the CLI

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

About this task

Copy Services features and RAID require that small amounts of volume cache be converted from cache memory into bitmap memory to allow the functions to operate. If you do not have enough bitmap space allocated when you try to use one of the functions, you will not be able to complete the configuration.

The total memory that can be dedicated to these functions is not defined by the physical memory in the system. The memory is constrained by the software functions that use the memory.

In planning the installation for a system, consider the future requirements for the advanced functions.

The following tables describe the amount of bitmap space necessary to configure the various copy services functions and RAID.

This table provides an example of the amount of memory that is required for remote mirroring functions, FlashCopy functions, and volume mirroring.

Table 10. Examples of memory required

Function	Grain size	1 MiB of memory provides the following volume capacity for the specified I/O group
Remote copy	256 KiB	2 TiB of total Metro Mirror, Global Mirror, or HyperSwap volume capacity
FlashCopy	256 KiB	2 TiB of total FlashCopy source volume capacity
FlashCopy	64 KiB	512 GiB of total FlashCopy source volume capacity
Incremental FlashCopy	256 KiB	1 TiB of total incremental FlashCopy source volume capacity
Incremental FlashCopy	64 KiB	256 GiB of total incremental FlashCopy source volume capacity
Volume mirroring	256 KiB	2 TiB of mirrored volume capacity

Table 10. Examples of memory required (continued)

Function	Grain size	1 MiB of memory provides the following volume capacity for the specified I/O group
Note: <ol style="list-style-type: none"> For multiple FlashCopy targets, you must consider the number of mappings. For example, for a mapping with a grain size of 256 KiB, 8 KiB of memory allows one mapping between a 16 GiB source volume and a 16 GiB target volume. Alternatively, for a mapping with a 256 KiB grain size, 8 KiB of memory allows two mappings between one 8 GiB source volume and two 8 GiB target volumes. When creating a FlashCopy mapping, if you specify an I/O group other than the I/O group of the source volume, the memory accounting goes toward the specified I/O group, not toward the I/O group of the source volume. For volume mirroring, the full 512 MiB of memory space enables 1 PiB of total volume mirroring capacity. When creating new FlashCopy relationships or mirrored volumes, additional bitmap space is allocated automatically by the system if required. 		

Table 11 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 11. RAID level comparisons

Level	Member count	Approximate capacity	Redundancy	Approximate bitmap memory cost
RAID-0	1-8	$MC * MS$	None	$(1 \text{ MB per } 2 \text{ TB of } MS) * MC$
RAID-1	2	MS	1	$(1 \text{ MB per } 2 \text{ TB of } MS) * (MC/2)$
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	$(1 \text{ MB per } 2 \text{ TB of } MS) * (MC/2)$
Note: There is a margin of error on the approximate bitmap memory cost of approximately 15%. For example, the cost for a 256 KB strip size for RAID-5 is ~1.15 MB for the first 2 TB of MS .				

Before you specify the configuration changes, consider the following factors.

- For FlashCopy mappings, only one I/O group consumes bitmap space. By default, the I/O group of the source volume is used.
- For Metro Mirror, Global Mirror, and HyperSwap active-active relationships, two bitmaps exist. For Metro Mirror or Global Mirror relationships, one is used for the master clustered system and one is used for the auxiliary system because the direction of the relationship can be reversed. For active-active relationships, which are configured automatically when HyperSwap volumes are created, one bitmap is used for the volume copy on each site because the direction of these relationships can be reversed.
- When you create a reverse mapping; for example, to run a restore operation from a snapshot to its source volume; a bitmap is also created for this reverse mapping.
- When you configure change volumes for use with Global Mirror or Metro Mirror, two internal FlashCopy mappings are created for each change volume.
- The smallest possible bitmap is 4 KiB; therefore, a 512 byte volume requires 4 KiB of bitmap space.

On existing systems, also consider these factors:

- When you create FlashCopy mappings and mirrored volumes, HyperSwap volumes, or formatted, fully allocated volumes, the system attempts to automatically increase the available bitmap space. You do not need to manually increase this space.
- Metro Mirror and Global Mirror relationships do not automatically increase the available bitmap space. You might need to use the **chiogrp** command or the management GUI to manually increase the space in one or both of the master and auxiliary systems.

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

Procedure

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

```
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:

```
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_grp0
node_count 2
vdisk_count 40
host_count 1
flash_copy_total_memory 5.0MB
flash_copy_free_memory 5.0MB
remote_copy_total_memory 20.0MB
remote_copy_free_memory 20.0MB
mirroring_total_memory 20.0MB
mirroring_free_memory 20.0MB
raid_total_memory 40.0MB
raid_free_memory 0.1MB
maintenance no
compression_active no
accessible_vdisk_count 40
compression_supported yes
max_enclosures 21
encryption_supported yes
```

Creating volumes using the CLI

You can use the command-line interface (CLI) to create a volume. You can create volumes that are not high availability volumes or you can create high availability volumes.

Before you begin

If the volume that you are creating maps to a flash drive, the data that is stored on the volume is not protected against Flash drive failures or node failures. To avoid data loss, add a volume copy that maps to an Flash drive on another node.

This task assumes that the clustered system (system) has been set up and that you have created storage pools. You can establish an empty storage pool to hold the MDisks that are used for image mode volumes.

About this task

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

Use the **mkvdisk** command to create sequential, striped, or image mode volumes that are not high availability volumes. Use the **mkvolume** command to create high availability volumes (or volumes that are not high availability). Use the **mkimagevolume** command to create an image mode volume by importing (preserving) data on a managed disk from another storage system.

Procedure

To create volumes, complete these steps.

1. Issue the **lsmdiskgrp** CLI command to list the available storage pools and the amount of free storage in each group.

Issue this CLI command to list storage pools:

```
lsmdiskgrp -delim :
```

This output 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 storage pool you want to provide the storage for the volume.
3. Issue the **lsiogrp** CLI command to show the I/O groups and the number of volumes assigned to each I/O group.

Note: It is normal for systems with more than one I/O group to have **mkvdisk** that have volumes in different I/O groups. You can use FlashCopy to make copies of volumes regardless of whether the source and target volume are in the same I/O group.

Similarly, if you plan to use intra-system Metro Mirror or Global Mirror, both the master and auxiliary volume can be in the same I/O group or different I/O groups.

Issue this CLI command to list I/O groups:

```
lsiogrp -delim :
```

This output 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 volume to. This determines which SAN Volume Controller nodes in the system process the I/O requests from the host systems. If you have more than one I/O group, make sure you distribute the volumes between the I/O groups so that the I/O workload is shared evenly between all SAN Volume Controller nodes.
5. Issue the **mkvdisk** CLI command to create a volume (that is not a high availability volume) that uses striped virtualization. Use the **mkvolume** command to create high availability volumes.

The rate at which the volume copies resynchronize after loss of synchronization can be specified by using the **-syncrate** parameter. Table 12 on page 30 defines the rates. These settings also affect the initial rate of formatting.

Table 12. 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
81-90	32 MB
91-100	64 MB

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

Issue this CLI command to create a volume with two copies using the I/O group and storage pool name and specifying the synchronization rate:

```
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 volume to use, *grpa* is the name of the storage pool for the primary copy of the volume and *grpb* is the name of the storage pool for the second copy of the volume, and 2 is the number of volume copies and the synchronization rate is 90 which is equivalent to 32MB per second.

Issue this CLI command to create a volume using the I/O group ID and storage pool ID:

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

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

Issue this CLI command to create a thin-provisioned volume using the I/O group and storage pool name:

```
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 volume to use and 20% is how much real storage to allocate to the volume, as a proportion of its virtual size. In this example, the size is 10 GB so that 2 GB will be allocated.

Issue this CLI command to create a volume with two copies using the I/O group and storage pool name:

```
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 volume to use, *grpa* is the name of the storage pool for the primary copy of the volume and *grpb* is the name of the storage pool for the second copy of the volume, and 2 is the number of volume copies.

Issue this CLI command to create a striped high availability volume:

```
mkvolume -pool 0:1 -size 1000
```

This creates a volume in storage pool 0 with a capacity of 1000 MBs.

Issue this CLI command to create an image mode volume:

```
mkimagevolume -mdisk 7 -pool 1 -thin -size 25 -unit gb
```


This imports a space-efficient image mode volume with a virtual capacity of 25 GB in storage pool 1 using MDisk 7.

Note: If you want to create two volume copies of different types, create the first copy using the **mkvdisk** command and then add the second copy using the **addvdiskcopy** command. To create a high availability volume, use the **mkvolume** command. To convert a basic volume to a high availability volume use the **addvolumecopy** command.

6. Issue the **lsvdisk** CLI command to list all the volumes that have been created.

Adding a copy to a volume

You can use the management GUI or command-line interface (CLI) to add a mirrored copy to a volume. Each volume can have a maximum of two copies.

Before you begin

The system supports mirrored copies for both standard topology, which consists of a single site, and HyperSwap and stretched system topologies, which consist of multiple sites. Both HyperSwap and stretched system topologies are used for high availability configurations. However, in HyperSwap topologies, separate I/O groups are at each site. For stretched system, individual I/O groups are split between sites with each node in the I/O group at separate sites. If you are adding a mirrored copy to a volume used in a standard topology, use the **addvdiskcopy** command or the management GUI. For HyperSwap and stretched volumes, use the **addvolumecopy**. In the management GUI, select **Volumes**, right-click the volume, and select **Add Volume Copies**.

Deleting a copy from a volume

You can use the management GUI or command-line interface (CLI) to delete a copy from a volume.

Before you begin

The system supports mirrored copies for both standard topology, which consists of a single site, and HyperSwap and stretched system topologies, which consist of multiple sites. Both HyperSwap and stretched system topologies are used for high availability configurations. However, in HyperSwap topologies, separate I/O groups are at each site. For stretched system, individual I/O groups are split between sites with each node in the I/O group at separate sites. The system supports mirrored copies for both standard topology, which consists of a single site, and HyperSwap and stretched system topologies, which consist of multiple sites. Both HyperSwap and stretched system topologies are used for high availability configurations. However, in HyperSwap topologies, separate I/O groups are at each site. For stretched system, individual I/O groups are split between sites with each node in the I/O group at separate sites. If you are deleting a copy to a volume used in a single system, use the **rmvdiskcopy** command or the management GUI. For HyperSwap volumes, use the **rmvolumecopy**. In the management GUI, select **Volumes**, right-click the volume copy, and select **Delete this Copy**.

Configuring host objects using the CLI

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

Before you begin

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 clustered system (system) that uses iSCSI connections, ensure that you have completed the necessary host-system configurations and have configured the system for iSCSI connections.

At least one WWPN or iSCSI name must be specified.

About this task

Use the following steps to create host objects:

Procedure

1. Issue the **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.

This is an example of the CLI command that you can issue to create a Fibre Channel attached host:

```
mkhost -name new_name -fcwwpn 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:

```
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 **addhostport** CLI command.

For example, issue the following CLI command:

```
addhostport -hbawwpn wwpn_list new_name
```

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

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

For example, issue the following CLI command:

```
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 **chhost** CLI command. This secret is shared between the host and the system. For example, issue the following CLI command:

```
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 **lsiscsiauth** command. To clear any previously set CHAP secret for a host, use the **chhost -nochapsecret** command.

What to do next

After you have created the host object on the system, 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 system. 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 **lsfabric** CLI command or the View Fabric panel under Service and Maintenance container in the management GUI.

Creating host mappings by using the CLI

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

About this task

Note: To understand the CLI for creating shared mappings in a host cluster, see the information about host clusters and the **mkhostcluster** command.

To create host mappings, follow these steps:

Procedure

1. Issue the **mkvdiskhostmap** CLI command to create host mappings.
This example is a CLI command that you can issue to create host mappings:

```
mkvdiskhostmap -host demohost1 mainvdisk1
```


Where *demohost1* is the name of the host and *mainvdisk1* is the name of the volume.
2. After you map volumes 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 system. You also have the option of creating a file system for those new disks. For more information about completing this task, see your host system documentation.

Creating FlashCopy mappings using the CLI

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

Before you begin

A FlashCopy mapping specifies the source and target volume. Source volumes and target volumes must meet these requirements:

- They must be the same size
- They must be managed by the same clustered system (system)

About this task

A volume can be the source in up to 256 mappings. A mapping is started at the point in time when the copy is required.

This task creates FlashCopy mappings:

Procedure

1. The source and target volume must be the exact same size. Issue the **lsvdisk -bytes** CLI command to find the size (capacity) of the volume in bytes.

2. Issue the **mkfcmap** CLI command to create a FlashCopy mapping.

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

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

Where *mainvdisk1* is the name of the source volume, *bkpvdisk1* is the name of the volume that you want to make the target volume, *main1copy* is the name that you want to call the FlashCopy mapping, and 75 is the copy rate (which translates to MB per second).

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

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

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

Note: The default copy rate of 50 (which translates to 2 MB per second) is used if you do not specify a copy rate.

If the specified source and target volumes are also the target and source 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 **lsfcmap** CLI command to check the attributes of the FlashCopy mappings that have been created:

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

```
lsfcmap -delim :
```

Where **-delim** specifies the delimiter. This 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.

About this task

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

These steps help you prepare and start a FlashCopy mapping:

Procedure

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

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

```
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 volume 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 **lsfcmap** CLI command to check the state of the mapping.

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

```
lsfcmap -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:main1disk1:1:bkpvdisk1:::prepared:0:50
```

3. Issue the **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:

```
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 volume is being used as a source in another active FlashCopy mapping.

4. Issue the **lsfcmapprogress** 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.

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

Results

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

Stopping FlashCopy mappings by using the CLI

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

About this task

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

Procedure

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

```
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 that is associated with the mapping and break the dependency on the source volume of any mappings that are also dependent on the target disk, issue the following command:

```
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.

Important: Using the **force** parameter might result in a loss of access. Use it only under the direction of the IBM Support Center.

The **split** parameter can be specified only when stopping a map that has a progress of 100 as shown by the **lsfcmap** command. The **split** parameter removes the dependency of any other mappings on the source volume. It might be used before starting another FlashCopy mapping whose target disk is the source disk of the mapping that is 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.

Before you begin

The **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 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 **lsfcmapdependentmaps** command.

About this task

Procedure

1. To delete an existing mapping, issue the **rmfcmap** CLI command:
`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 volume online, issue the following command:
`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.

Results

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.

About this task

If you have created several FlashCopy mappings for a group of 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.

Procedure

To add FlashCopy mappings to a new FlashCopy consistency group, complete the following steps.

1. Issue the **mkfcconsistgrp** CLI command to create a FlashCopy consistency group.
The following CLI command is an example of the command you can issue to create a FlashCopy consistency group:
`mkfcconsistgrp -name FCgrp0 -autodelete`
Where `FCgrp0` 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 **lsfcconsistgrp** CLI command to display the attributes of the group that you have created.

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

```
lsfcconsistgrp -delim : FCcgrp0
```

The following output 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 **chfcmap** CLI command to add FlashCopy mappings to the FlashCopy consistency group:

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

```
chfcmap -consistgrp FCcgrp0 main1copy
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 **lsfcmap** CLI command to display the new attributes of the FlashCopy mappings.

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

```
lsfcmap -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 **lsfcconsistgrp** CLI command to display the detailed attributes of the group.

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

```
lsfcconsistgrp -delim : FCcgrp0
```

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

The following output 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.

About this task

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

issued to prepare every FlashCopy mapping in the group; only a single start command is issued to start every FlashCopy mapping in the group.

Procedure

To prepare and start a FlashCopy consistency group, complete the following steps.

1. Issue the **prestartfcconsistgrp** CLI command to prepare the FlashCopy consistency group. This command must be issued before the copy process can begin.

Remember: A single prepare command prepares all of the mappings simultaneously for the entire group.

An example of the CLI command issued to prepare the FlashCopy consistency group:

```
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 volume is being used as a source volume in 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 **lsfcconsistgrp** command to check the status of the FlashCopy consistency group.

An example of the CLI command issued to check the status of the FlashCopy consistency group:

```
lsfcconsistgrp -delim :
```

An example of the output displayed:

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

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

Remember: A single start command starts all the mappings simultaneously for the entire group.

An example of the CLI command issued to start the FlashCopy consistency group mappings:

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

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

Include the **prep** parameter, and the system automatically issues the **prestartfcconsistgrp** command for the specified group.

Note: Combining the **restore** parameter with the **prep** parameter, force-starts the consistency group. This occurs even if the target volume is being used as a source volume 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 **lsfcconsistgrp** command to check the status of the FlashCopy consistency group.

An example of the CLI command issued to check the status of the FlashCopy consistency group:

```
lsfcconsistgrp -delim : maintobkpfcopy
```

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

An example of the output displayed during the copying process:

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

An example of the output displayed when the process copying is complete:


```
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.

Before you begin

The **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.

About this task

Procedure

1. To stop a FlashCopy consistency group, issue the **stopfcconsistgrp** CLI command:
`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 volumes of any mappings that are also dependent on the target volume, issue the following command:

```
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 volumes. 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

Results

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.

Before you begin

The **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.

About this task

Follow these steps to delete an existing consistency group:

Procedure

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

```
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.

Important: Using the **-force** parameter might result in a loss of access. Use it only under the direction of your support center.

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 **rmfcmap** command.

Results

The command does not return any output.

Creating Metro Mirror, Global Mirror, or active-active relationships using the CLI

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

About this task

These steps help create Metro Mirror, Global Mirror, or active-active relationships:

Procedure

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

```
mkrcrelationship -master master_volume_id  
-aux aux_volume_id -cluster system_id
```

Where *master_volume_id* is the ID of the master volume, *aux_volume_id* is the ID of the auxiliary volume, and *system_id* is the ID of the remote clustered system.

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

```
mkrcrelationship -master master_volume_id  
-aux aux_volume_id -cluster system_id -global
```

Where *master_volume_id* is the ID of the master volume, *aux_volume_id* is the ID of the auxiliary volume, and *system_id* is the ID of the remote system.

3. To create a new relationship with cycling enabled:

```
mkrcrelationship -master books_volume -aux books_volume -cluster DR_cluster -global -cyclingmode multi
```

Note: Add change volumes to a relationship by issuing **chrcrelationship -auxchange** or **chrcrelationship -masterchange**.

4. To create a new active-active relationship, run the **mkrcrelationship** command with the **-activeactive** parameter. For example, enter:

```
mkrcrelationship -master master_volume_id -aux aux_volume_id -cluster system_id -activeactive
```

Where *master_volume_id* is the ID of the master volume, *aux_volume_id* is the ID of the auxiliary volume, and *system_id* is the ID of the remote system.

Modifying Metro Mirror, Global Mirror, or active-active relationships using the CLI

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

About this task

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

Procedure

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

```
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.

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

```
chrcrelationship -force -noconsistgrp rc_rel_name/id
```

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

Important: Using the **-force** parameter might result in a loss of access. Use it only under the direction of your support center.

Starting and stopping Metro Mirror, Global Mirror, or active-active relationships using the CLI

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

About this task

Complete these steps to start or stop Metro Mirror, Global Mirror or active-active relationships:

Procedure

1. To start a Metro Mirror, Global Mirror, or an active-active relationship, run the **startcrrelationship** command. For example, enter:

```
startcrrelationship rc_rel_id
```

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

Note: Active-active relationships can only be started if they have a state of idling.

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

For example, enter:

```
stopcrrelationship rc_rel_id
```

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

3. To stop an active-active relationship, the following conditions must be met:

- The **-access** parameter is specified.
- The state of the relationship is **consistent_copying**.
- The status of the relationship is **primary_offline**.

For example, enter:

```
stopcrrelationship rc_rel_id -access
```

Where *rc_rel_id* is the ID of the active-active relationship that you want to stop. The **-access** parameter allows hosts read or write access to a volume in an active-active relationship that contains an older but a consistent image that can be used in a disaster recovery scenario.

Displaying the progress of Metro Mirror, Global Mirror, or active-active relationships using the CLI

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

About this task

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

Procedure

1. To display data progress without headings for columns of data or for each item of data in a Metro Mirror, Global Mirror, or active-active relationship, run the **lsrrelationshipprogress -nohdr** command. For example, to display data of the relationship with headings suppressed, enter the following command, where *rc_rel_name* is the name of the specified object type.

```
lsrrelationshipprogress -nohdr rc_rel_name
```

2. To display the progress of a background copy of a Metro Mirror, Global Mirror, or active-active 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:

```
lsrrelationshipprogress -delim :
```

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 volumes in a stand-alone Metro Mirror or Global Mirror relationship when that relationship is in a consistent state. You cannot switch roles between primary and secondary volumes for an active-active relationships.

About this task

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 volumes in Metro Mirror or Global Mirror relationships, follow these steps:

Procedure

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

```
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 **switchrcrelationship -primary aux** command. For example, enter:

```
switchrcrelationship -primary aux rc_rel_id
```

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

Remember:

- You cannot switch a global relationship if cycling is (automatically) set.

- To switch the direction of a **multi** cycling mode-based relationship, the relationship must stop with access enabled. Then, start by using **-force** in the opposite direction. (Using the **force** parameter might result in a loss of access. Use it only under the direction of your support center.)

Deleting Metro Mirror, Global Mirror, or active-active relationships using the CLI

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

Procedure

To delete Metro Mirror, Global Mirror, or active-active relationships, run the **rmrcrelationship** command. For example, enter:

```
rmrcrelationship rc_rel_name/id
```

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

Creating Metro Mirror, Global Mirror, or active-active consistency groups using the CLI

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

About this task

To create Metro Mirror, Global Mirror, or active-active consistency groups, complete these steps:

Procedure

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

```
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, Global Mirror, or active-active relationships to the group, run the **chrcrelationship** command. For example, enter:

```
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, Global Mirror, or active-active 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, Global Mirror, or active-active consistency group.

About this task

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

Procedure

1. Run the **chrcconsistgrp** command to assign a new name to the consistency group. For example, enter:
`chrcconsistgrp -name new_name_arg`
where *new_name_arg* is the assigned new name of the consistency group.
2. Run the **chrcconsistgrp** command to change the name of the consistency group. For example, enter:
`chrcconsistgrp -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, Global Mirror, or active-active consistency-group copy processes using the CLI

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

About this task

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

Procedure

1. To start a Metro Mirror, Global Mirror, or active-active consistency-group copy process, set the direction of copy if it is undefined and optionally mark the secondary volumes of the consistency group as clean. Run the **startrcconsistgrp** command. For example, enter:
`startrcconsistgrp rc_consist_group_id`
where *rc_consist_group_id* is the ID of the consistency group to start processing.

Note: If you are starting an active-active consistency group, all relationships in the group must be in **idling** state for the consistency group to start.

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

For example, enter:

```
stoprcconsistgrp 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 volumes in the group.

3. To stop the copy process for an active-active consistency group, the following conditions must be met:
 - The **-access** parameter is specified.
 - The state of the relationships in the consistency group are **consistent_copying**.
 - The status of the relationships in the consistency group are **primary_offline**.

For example, enter:

```
stoprcconsistgrp rc_consist_group_id -access
```

Where *rc_rel_id* is the ID of the active-active consistency group that you want to stop. The **-access** parameter allows hosts read or write access to a volume in an active-active relationship that contains an older but a consistent image that can be used in a disaster recovery scenario.

Deleting Metro Mirror, Global Mirror, or active-active consistency groups using the CLI

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

About this task

To delete existing Metro Mirror, Global Mirror, or active-active consistency groups, complete these steps:

Procedure

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

```
rmrconsistgrp rc_consist_group_id
```

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

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

```
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.

Important: Using the **force** parameter might result in a loss of access. Use it only under the direction of your support center.

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.

Procedure

Note: When remote copy partnerships are created between systems supporting different maximum numbers of volumes, then the maximum number of volumes that can be created on any system is determined to be the same as on the system that supports the smallest maximum number of volumes. An attempt to create a partnership will fail if one system has more disks than is supported by the other system.

To create Metro Mirror and Global Mirror partnerships, complete the following steps.

1. To create Metro Mirror and Global Mirror partnerships for Fibre Channel connections, run the **mkfcpartnership** command. To create Metro Mirror and Global Mirror partnerships for IP connections, run the **mkippartnership** command. For example, for Fibre Channel connections enter:

```
mkfcpartnership -linkbandwidthmbits bandwidth_in_mbps  
-backgroundcopyrate percentage_of_available_bandwidth 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, *percentage_of_available_bandwidth* specifies the maximum percentage of aggregate link bandwidth that can be used for background copy operations, and *remote_cluster_id* is the ID of the remote system. For IP connections, enter:

```
mkippartnership -type ip_address_type  
-clusterip remote_cluster_ip_address  
-chapsecret chap_secret  
-linkbandwidthmbits bandwidth_in_mbps  
-backgroundcopyrate percentage_of_available_bandwidth
```

where *ip_address_type* specifies the IP address type (IPv4 or IPv6) that is used by the background copy process between the clusters, *remote_cluster_ip_address* specifies the IP address of the remote system, *chap_secret* specifies the CHAP secret of the remote system (optional), *bandwidth_in_mbps* specifies the bandwidth (in megabytes per second) that is used by the background copy process between the clusters, and *percentage_of_available_bandwidth* specifies the maximum percentage of aggregate link bandwidth that can be used for background copy operations (optional).

2. Run the **mkfcpartnership** command for Fibre Channel connections or **mkippartnership** command for IP connections from the remote system. For example, for Fibre Channel connections enter:

```
mkfcpartnership -linkbandwidthmbits bandwidth_in_mbps
                 -backgroundcopyrate percentage_of_available_bandwidth
                 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, *percentage_of_available_bandwidth* specifies the maximum percentage of aggregate link bandwidth that can be used for background copy operations, and *remote_cluster_id* is the ID of the local system.

For Internet Protocol (IP) connections, enter:

```
mkippartnership -type ip_address_type
                 -clusterip local_cluster_ip_address
                 -chapsecret chap_secret
                 -linkbandwidthmbits bandwidth_in_mbps
                 -backgroundcopyrate percentage_of_available_bandwidth
```

where *ip_address_type* specifies the IP address type (IPv4 or IPv6) that is used by the background copy process between the clusters, *remote_cluster_ip_address* specifies the IP address of the local system, *chap_secret* specifies the CHAP secret of the local system (optional), *bandwidth_in_mbps* specifies the bandwidth (in megabytes per second) that is used by the background copy process between the clusters, and *percentage_of_available_bandwidth* specifies the maximum percentage of aggregate link bandwidth that can be used for background copy operations (optional).

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.

About this task

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

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

Procedure

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

```
chpartnership -type ip_address_type
               -clusterip remote_cluster_ip_address
               -chapsecret chap_secret
               -nochapsecret -linkbandwidthmbits bandwidth_in_mbps
               -backgroundcopyrate percentage_of_available_bandwidth remote_cluster_id
```

where *ip_address_type* specifies the IP address type ("ipv4" or "ipv6") that is used by the background copy process between the clusters (only used for IP connections), *remote_cluster_ip_address* specifies the IP address of the remote cluster (only used for IP connections), *chap_secret* specifies the CHAP secret of the remote cluster (only used for IP connections), *bandwidth_in_mbps* specifies the bandwidth (in megabytes per second) that is used by the background copy process between the clusters (this is optional), *percentage_of_available_bandwidth* specifies the maximum percentage of aggregate link bandwidth that can be used for background copy operations (this is optional), and *remote_cluster_id* is the ID or name of the remote system.

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

```
chpartnership -type ip_address_type
               -clusterip local_cluster_ip_address
               -chapsecret chap_secret -nochapsecret
               -linkbandwidthmbits bandwidth_in_mbps
               -backgroundcopyrate percentage_of_available_bandwidth local_cluster_id
```

where *ip_address_type* specifies the IP address type ("ipv4" or "ipv6") that is used by the background copy process between the clusters (only used for IP connections), *local_cluster_ip_address* specifies the IP address of the local cluster (only used for IP connections), *chap_secret* specifies the CHAP secret of

the local cluster (only used for IP connections), *bandwidth_in_mbps* specifies the bandwidth (in megabytes per second) that is used by the background copy process between the clusters (this is optional), *percentage_of_available_bandwidth* specifies the maximum percentage of aggregate link bandwidth that can be used for background copy operations (this is optional), and *local_cluster_id* is the ID or name of the local system.

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.

About this task

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

Procedure

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

```
chpartnership -start cluster_id
```

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

```
chpartnership -stop cluster_id
```

Where *cluster_id* is the ID of the local or 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.

About this task

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

Procedure

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:

```
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 **rmpartnership** command from either cluster. For example, enter:

```
rmpartnership remote_cluster_id
```

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

Determining the WWNNs of a node using the CLI

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

About this task

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

Procedure

1. Issue the **lsnode** CLI command to list the nodes in the clustered system.
2. Record the name or ID of the node for which you want to determine the WWNNs.
3. Issue the **lsportfc** 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:

```
lsportfc -filtervalue node_id=2
```

Where `node_id=2` is the name of the node for which you want to determine the WWNNs. The following is the output from the command:

id	fc_io_port_id	port_id	type	port_speed	node_id	node_name	WWNN	nportid	status
0	1	1	fc	8Gb	2	node2	5005076801405F82	010E00	active
1	2	2	fc	8Gb	2	node2	5005076801305F82	010A00	active
2	3	3	fc	8Gb	2	node2	5005076801105F82	010E00	active
3	4	4	fc	8Gb	2	node2	5005076801205F82	10A00	active
4	5	3	ethernet	10Gb	2	node2	5005076801505F82	540531	active
5	6	4	ethernet	10Gb	2	node2	5005076801605F82	E80326	active

4. Record the six WWNNs (to assist with setting up other systems).

Listing node-dependent volumes using the CLI

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

Before you begin

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

About this task

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 flash drives 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 Flash drive MDisk to complete. To resolve (2-4), bring any offline MDisk online and repair any degraded paths.

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

Procedure

1. Issue the **lsdependentvdisks** CLI command.

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

```
lsdependentvdisks -enclosure -delim : 0:1
```

The following example shows the output that is displayed:

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

2. If the **lsdependentvdisks** 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 **lsdependentvdisks** command. When the command returns no volumes, the system is free from any node-dependent volumes.

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

```
lsdependentvdisks -node01 :
```

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

```
vdisk_id      vdisk_name
```

Determining the volume name from the device identifier on the host

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

About this task

Each volume that is exported by the system is assigned a unique device identifier. The device identifier uniquely identifies the volume and can be used to determine which volume corresponds to the volume that the host detects.

Complete the following steps to determine the volume name from the device identifier:

Procedure

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 system 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 must go into the HBA BIOS.
 - b. Verify which host object is defined to the system 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:

```
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 host mappings:

```
lshostvdiskmap hostname
```

Where *hostname* is the name of the host.

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

Determining the host that a volume maps

You can determine the host that a volume maps by using the command-line interface (CLI). To view the host mapping for a volume in the management GUI, select **Volumes > Volumes by Hosts**.

About this task

Complete the following steps to determine the host that the volume maps:

Procedure

1. Enter the following CLI command to list the hosts to which this volume maps:
`lsvdiskhostmap vdisk_name | vdisk_id`
where *vdisk_name* | *vdisk_id* is the name or ID of the volume.
2. Find the host name or ID to determine which host this volume maps.
 - If no data is returned, the volume does not map any hosts.

Determining the relationship between volume and MDisks using the CLI

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

About this task

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

Procedure

- To display a list of the IDs that correspond to the MDisks that comprise the volume, issue the following CLI command:
`lsvdiskmember vdiskname/id`
where *vdiskname/id* is the name or ID of the volume.
- To display a list of IDs that correspond to the volumes that are using this MDisk, issue the following CLI command:
`lsmdiskmember mdiskname/id`
where *mdiskname/id* is the name or ID of the MDisk.
- To display a table of volume IDs and the corresponding number of extents that are being used by each volume, issue the following CLI command:
`lsmdiskextent 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 volume, issue the following CLI command:
`lsvdiskextent vdiskname/id`
where *vdiskname/id* is the name or ID of the volume.

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).

About this task

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.

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

Procedure

1. Enter the following command to display a detailed view of the MDisk:
`lsmdisk object_name`
Where *object_name* 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. Enter the following command to display a detailed view of the controller:
`lscontroller controller_name`
Where *controller_name* 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 system by using the CLI

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

About this task

Complete the following steps to increase the size of your system:

Procedure

1. Add a node to your system and repeat this step for the second node.
2. Migrate your volumes to new I/O groups if you want to balance the load between the existing I/O groups and the new I/O groups. Repeat this step for all volumes that you want to assign to the new I/O group.

Adding a node to increase the size of the system

You can add a node to the system by using the CLI or management GUI. A node can be added to the system if the node previously failed and is being replaced with a new node or if a repair action has caused the node to be unrecognizable by the system. When you add nodes, ensure that they are added in pairs to create a full I/O group. Adding a node to the system increases the capacity of the entire system.

You can use either the management GUI or the command-line interface to add a node to the system. Some models might require using the front panel to verify that the new node was added correctly.

Before you add a node to a system, you must make sure that the switch zoning is configured such that the node that is being added is in the same zone as all other nodes in the system. 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 that is being added is in the same VSAN or zone.

Considerations when adding a node to a system

If you are adding a node that was used previously, either within a different I/O group within this system or within a different system, take into account that if you add a node without changing its worldwide node name (WWNN), hosts might detect the node and use it as if it were in its old location. This action might cause the hosts to access the wrong volumes.

- You must ensure that the model type of the new node is supported by the software level that is installed on the system. If the model type is not supported by the software level, update the system to a software level that supports the model type of the new node.
- Each node in an I/O group must be connected to a different uninterruptible power supply.
- If you are adding a node back to the same I/O group after a service action required it to be deleted from the system, and if the physical node has not changed, then no special procedures are required to add it back to the system.
- If you are replacing a node in a system either because of a node failure or an update, you must change the WWNN of the new node to match that of the original node before you connect the node to the Fibre Channel network and add the node to the system.
- If you are adding a node to the SAN again, to avoid data corruption, ensure that you are adding the node to the same I/O group from which it was removed. You must use the information that was recorded when the node was originally added to the system. If you do not have access to this information, contact the support center for assistance with adding the node back into the system so there is no data corruption.
- For each external storage system, 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 system. You must ensure that the LUNs are the same before you add the new node to the system.
- If you are creating an I/O group in the system and are adding a new node, there are no special procedures because this node was never added to a system and the WWNN for the node did not exist.
- If you are creating an I/O group in the system and are adding a new node that was added to a system before, the host system might still be configured to the node WWPNs and the node might still be zoned in the fabric. Because you cannot change the WWNN for the node, you must ensure that other components in your fabric are configured correctly. Verify that any host that was previously configured to use the node was correctly updated.
- If the node that you are adding was previously replaced, either for a node repair or update, you might have used the WWNN of that node for the replacement node. Ensure that the WWNN of this node was updated so that you do not have two nodes with the same WWNN attached to your fabric. Also, ensure that the WWNN of the node that you are adding is not 00000. If it is 00000, contact your support representative.
- The new node must be running a software level that supports encryption.
- If you are adding the new node to a system with either a HyperSwap or stretched system topology, you must assign the node to a specific site.

Considerations when using multipathing device drivers

- 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 that are supported by the multipathing device drivers. Multipathing device drivers maintain an association between a vpath and a volume. This association uses an identifier (UID) which is unique to the volume and is never reused. The UID allows multipathing device drivers to directly associate vpaths with volumes.
- Multipathing device drivers operate within a protocol stack that contains disk and Fibre Channel device drivers that are used to communicate with the system by using the SCSI protocol over Fibre Channel as defined by the ANSI FCS standard. The addressing scheme that is 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 by using the same WWNN and LUN numbers that were previously used.
- Multipathing device drivers do not check the association of the volume with the vpath on every I/O operation that it performs.

After the new node is zoned and cabled correctly to the existing system, you can use either the **addnode** command or the **Add Node** wizard in the management GUI. To access the **Add Node** wizard, select **Monitoring > System**. On the image, click the new node to launch the wizard. Complete the wizard and verify the new node. If the new node is not displayed in the image, it indicates a potential cabling issue. Check the installation information to ensure that your node was cabled correctly.

To add a node to a system by using the command-line interface, complete these steps:

1. Enter this command to verify that the node is detected on the fabric:

```
svcinfo lsnodecandidate
```

This example shows the output for this command:

```
# svcinfo lsnodecandidate
id                panel_name UPS_serial_number UPS_unique_id hardware serial_number product_mtm machine_signature
500507680C007B00 KD0N8AM          500507680C007B00          KD0N8AM          2145-DH8 0123-4567-89AB-CDEF
```

The **id** parameter displays the WWNN for the node. If the node is not detected, verify cabling to the node.

2. Enter this command to determine the I/O group where the node should be added:

```
lsiogrp
```

3. Record the name or ID of the first I/O group that has a node count of zero (0). You will need the name or 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.
4. Enter this command to add the node to the system:

```
addnode -wwnodename WWNN -iogrp iogrp_name -name new_name_arg -site site_name
```

Where **WWNN** is the WWNN of the node, **iogrp_name** is the name of the I/O group that you want to add the node to and **new_name_arg** 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 that you specify a meaningful name. The **site_name** specifies the name of the site location of the new node. This parameter is only required if the topology is a HyperSwap or stretched system.

Note: Adding the node might take a considerable amount of time.

5. Record this information for future reference:

- Serial number.
- Worldwide node name.
- All of the worldwide port names.
- The name or ID of the I/O group

To add either a SAN Volume Controller 2145-CG8 or SAN Volume Controller 2145-CF8 node to a system by using the command-line interface, complete these steps:

1. Use the front panel of the node, record the WWNN. The front panel only shows the last 5 digits of the WWNN.
2. Enter this command to verify that the node is detected on the fabric:

```
svcinfo lsnodecandidate
```

This example shows the output for this command:

```
# svcinfo lsnodecandidate
id          panel_name UPS_serial_number UPS_unique_id hardware serial_number product_mtm machine_signature
500507680100E85F 168167 UPS_Fake_SN 100000000000E85F C68 7860123 2145-DH8 0123-4567-89AB-CDEFsvcinfo lsnodecandidate
```

The **id** parameter displays the WWNN for the node. Ensure that the last 5 digits that are displayed match the WWNN on the front panel. If the node is not detected, verify cabling to the node.

3. Enter this command to determine the I/O group where the node should be added:

```
lsiogrp
```

4. 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.
5. Enter this command to add the node to the system:

```
addnode -wwnodename WWNN -iogrp iogrp_name -name newnodename -site newsitename
```

Where **WWNN** is the WWNN of the node, **iogrp_name** 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 that you specify a meaningful name. The **newsitename** specifies the name of the site location of the new node. This parameter is only required if the topology is a HyperSwap or stretched system.

Note: Adding the node might take a considerable amount of time.

6. Record this information for future reference:

- Serial number.
- Worldwide node name.
- All of the worldwide port names.
- The name or ID of the I/O group

If a node shows node error 578 or node error 690, the node is in service state. Complete the following steps from the front panel to exit service state:

1. Press and release the up or down button until the **Actions?** option is displayed.
2. Press the **Select** button.
3. Press and release the **Up** or **Down** button until the **Exit Service?** option is displayed.
4. Press the **Select** button.
5. Press and release the **Left** or **Right** button until the **Confirm Exit?** option is displayed.
6. Press the **Select** button.

Validating and repairing mirrored volume copies by using the CLI

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

Attention: Run the **repairvdiskcopy** command only if all 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 volume 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 volume 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 volume.

-medium

Use this parameter to convert sectors on all volume 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 volume copy to the other volume copy. The command corrects any differing sectors by copying the sectors from the primary copy to the copies that are 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 volume 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** by using this parameter to list all the differences.

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

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

Notes:

1. Only one **repairvdiskcopy** command can run on a volume at a time.
2. After you start the **repairvdiskcopy** command, you cannot use the command to stop processing.
3. The primary copy of a mirrored volume 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 that is 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 volume copies by using the CLI

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

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

```
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 thin-provisioned volume using the CLI

You can use the **repairsevdiskcopy** command from the command-line interface to repair the metadata on a thin-provisioned volume.

The **repairsevdiskcopy** command automatically detects and repairs corrupted metadata. The command holds the 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 volume copies in the configuration.

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

Attention: Use this command only to repair a thin-provisioned volume that has reported corrupt metadata.

Issue the following command to repair the metadata on a thin-provisioned volume:

```
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 thin-provisioned volume by using the CLI

Issue the **lsrepairsevdiskcopyprogress** command to list the repair progress for thin-provisioned volume copies of the specified volume. If you do not specify a volume, the command lists the repair progress for all thin-provisioned copies in the system.

Note: Only run this command after you run the **repairsevdiskcopy** command, which you must only run as required by the fix procedures recommended by your support team.

Recovering offline volumes using the CLI

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

About this task

If you lose both nodes in an I/O group, you will lose access to all volumes that are associated with the I/O group. To regain access to the volumes, you must perform one of the following procedures.. 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 flushed to the

backend. 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 system.

Complete the following steps to recover an offline volume when one node has down-level hardened data and the other node has lost hardened data.

Procedure

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

Example

Data loss scenario 2

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

Complete the following steps to recover an offline volume when both nodes that have lost their hardened data and cannot be recognized by the system.

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

Recovering a node and returning it to the system by 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 system.

About this task

Complete the following steps to recover a node and return it to the system.

Procedure

1. Run the **lsnode** command to verify that the node is offline.
2. Run the **rmnode** *nodename_or_ID* command to remove the old instance of the offline node from the system.
3. Run the **lsnodecandidate** command to verify that the node is visible on the fabric.
4. Run the **addnode** command to add the node back into the system. In the following command, *wwnn* is the worldwide node name, *iogroupname_or_ID* identifies the I/O group, and *nodename* is the name of the node.

```
addnode -wwnodename wwnn -iogrp iogroupname_or_ID -name nodename
```

Note: In a service situation, a node must be added back into a system that uses the original node name. If the partner node in the I/O group has not also been deleted, this is the default name that is used if the **-name** parameter is not specified.

5. Run the **lsnode** command to verify that the node is online.

Recovering offline volumes using the CLI

You can recover offline volumes using the command-line interface (CLI).

About this task

Perform the following steps to recover offline volumes:

Procedure

1. Issue the following CLI command to list all volumes that are offline and belong to an I/O group, enter:

```
lsvdisk -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 volume with a *fast_write_state* of **corrupt** and bring the volume back online, enter:

```
recovervdisk vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the volume.

Notes:

- If the specified volume is space-efficient or has space-efficient copies, the **recovervdisk** command starts the space-efficient repair process.
 - If the specified volume 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:

```
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 volume is space-efficient or has space-efficient copies, the **recovervdiskbyiogrp** command starts the space-efficient repair process.
 - If any volume is mirrored, the **recovervdiskbyiogrp** command starts the resynchronization process.
4. To acknowledge data loss for all volumes in the clustered system with a *fast_write_state* of **corrupt** and bring them back online, enter:

```
recovervdiskbycluster
```

Notes:

- If any volume is space-efficient or has space-efficient copies, the **recovervdiskbycluster** command starts the space-efficient repair process.
- If any volume is mirrored, the **recovervdiskbycluster** command starts the resynchronization process.

Moving offline volumes to their original I/O group using the CLI

You can move offline volumes to their original I/O group using the command-line interface (CLI).

About this task

After a node or an I/O group fails, you can use the following procedure to move offline volumes to their original I/O group. The system disables moving a volume if the selected volume is currently formatting. After the formatting completes, you can move the volume.

Attention: Do not move volumes to an offline I/O group. Ensure that the I/O group is online before you move the volume back to avoid any further data loss.

Complete the following steps to move offline volumes to their original I/O group:

Procedure

1. Enter the following command to move the volume back into the original I/O group.
In the example, *7* is the name of the node that you want to move the volume, *IOGRP3* identifies the I/O group that you want to migrate the volume to, and *DB_volume* identifies the volume that you want to migrate.

```
movevdisk -iogrp IOGRP3 -node 7 DB_volume
```
2. Enter the following command, where *IO_grpname_or_ID* is the name or ID of the original I/O group, to verify that the volumes are now online.

```
lsvdisk -filtervalue IO_group_name= IO_grpname_or_ID
```

Recording WWPN changes of replaced host HBAs

You can use the command-line interface (CLI) to record a change to a defined host object.

Before you begin

Because it is sometimes necessary to replace the host-bus adapter (HBA) that connects the host to the SAN, you must inform the system of the new worldwide port names (WWPNs) that this HBA contains.

Ensure that your switch is zoned correctly.

Procedure

To inform the system of a change to a defined host object, complete the following steps.

1. Issue this CLI command to list the candidate HBA ports:

```
lsfcportcandidate
```


or

```
lssasportcandidate
```


You will see a list of the HBA ports that are available for addition to host objects. One or more of these HBA ports will 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:

```
lshost
```
3. Issue the following CLI command to list the WWPNs that are currently assigned to the host object:

```
lshost 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:

```
addhostport -fcwwpn one or more existing port names  
separated by : hostobjectname/ID
```


where *one or more existing port names 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:

```
rmhostport -fcwwpn one or more existing port names  
separated by : hostobjectname/ID
```

where one or more existing WWPNs separated by a colon (:) are the WWPNs that are currently assigned to the host object and *hostobjectname/ID* is the name or ID of the host object.

Note: When volume protection is enabled for the system, the deletion of the specified host port fails if the following conditions are met:

- It is the last active port on the host.
- It is mapped to any volume that has received I/O within the specified volume protection interval.

If volume protection is enabled, and the host port being deleted is the last port for a host which is mapped to any volume which has received I/O within the defined volume protection time period, then the command will be failed. If multiple hosts are mapped to the same active volume, the system deletes the port if the host is offline.

Results

Any mappings that exist between the host object and the volumes are automatically applied to the new WWPNs. Therefore, the host sees the volumes as the same SCSI LUNs as before.

What to do next

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 volumes using the CLI

You can use the command-line interface (CLI) to expand a volume on Windows, AIX, or Linux systems.

About this task

Volumes that are mapped for FlashCopy cannot be expanded. The system disables expanding a volume if the selected volume is performing quick initialization. After the quick initialization completes, you can expand the volume.

Run Windows Update and apply all recommended updates to your system before you attempt to expand a volume that is mapped to a Windows host.

Determine the exact size of the source or master volume by issuing the following CLI command:

```
lsvdisk -bytes vdiskname
```

where *vdiskname* is the name of the volume for which you want to determine the exact size.

Volumes can be expanded under Windows concurrently with I/O operations.

You can expand volumes for the following reasons:

- To increase the available capacity on a particular volume that is already mapped to a host.
- To increase the size of a volume so that it matches the size of the source or master volume, and so that it can be used in a FlashCopy mapping or Metro Mirror relationship.

You cannot expand the capacity of any volume in a Global Mirror with change volumes relationship or in a HyperSwap relationship.

You can expand the capacity of volumes in Metro Mirror and Global Mirror relationships that are in `consistent_synchronized` state if those volumes are using thin-provisioned or compressed copies. You cannot expand the following types of volumes:

- Volumes in HyperSwap relationships or in Global Mirror relationships that are operating in cycling mode
- Volumes in relationships where a change volume is configured
- Volumes that have a fully allocated copy

You must expand both volumes in a relationship to maintain full operation of the system. Expand the secondary volume by the required capacity, and then expand the primary volume.

A volume that is not mapped to any hosts and does not contain any data can be expanded at any time. If the volume contains data that is in use, you can expand the volume if your host has a supported AIX or Microsoft Windows operating system.

See the software restrictions page on the following website for supporting information and restrictions on expanding volumes: www.ibm.com/support

Expanding a volume that is mapped to an AIX host

The SAN Volume Controller supports expanding the size of a volume if the AIX host is using AIX version 5.2 or later.

About this task

The AIX **chvg** command option can be used to expand the size of a physical volume that the Logical Volume Manager (LVM) uses. The physical volume can be expanded without interruptions to the use or availability of the system. For more information, see the *AIX System Management Guide Operating System and Devices*.

Expanding a volume that is mapped to a Microsoft Windows host by using the CLI

You can use the command-line interface (CLI) to expand the size of a volume that is mapped to a Microsoft Windows host.

About this task

Complete the following steps to expand a volume that is mapped to a Windows host:

Procedure

1. Issue the following CLI command to expand the volume:

```
expandvdisksize -size disk_size -unit data_unit vdisk_name/vdisk_id
```

 - *disk_size* is the capacity by which you want to expand the volume
 - *b | kb | mb | gb | tb | pb* is the *data_unit* to use with the capacity
 - *vdisk_name/vdisk_id* is the name of the volume or the ID of the volume to expand
2. On the Windows host, start the Computer Management application and open the Disk Management window under the Storage branch.

Results

You see the volume 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.

What to do next

If the Computer Management application was open before you expanded the volume, 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 volume using the CLI

You can reduce the size of a compressed or uncompressed volume by using the command-line interface (CLI).

About this task

Volumes can be reduced in size, if it is necessary. You can make a target or auxiliary volume the same size as the source or master volume when you create FlashCopy mappings, Metro Mirror relationships, or Global Mirror relationships. However, if the volume contains data, do not shrink the size of the disk. The system disables shrinking a volume if the selected volume is performing quick initialization. After the quick initialization completes, you can shrink the volume.

Attention:

1. It is difficult to anticipate how an operating system or file system uses the capacity in a volume. When you shrink a volume, capacity is removed from the end of the disk, whether or not that capacity is in use. Even if a volume has free capacity, do not assume that only unused capacity is removed when you shrink a volume.
2. If the volume contains data that is being used, *do not attempt under any circumstances to shrink a volume 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 **shrinkvdiskspace** command to shrink the physical capacity that is allocated to the particular volume by the specified amount. You can also shrink the virtual capacity of a thin-provisioned volume without altering the physical capacity that is assigned to the volume.

You cannot shrink the capacity of any volume in a Global Mirror with change volumes relationship or in a HyperSwap relationship.

You can shrink the capacity of volumes in Metro Mirror and Global Mirror relationships that are in `consistent_synchronized` state if those volumes are using thin-provisioned or compressed copies. You cannot shrink the following types of volumes:

- Volumes in HyperSwap relationships or in Global Mirror relationships that are operating in `cycling` mode.
- Volumes in relationships where a change volume is configured.
- Volumes that have a fully allocated copy.

You must shrink both volumes in a relationship to maintain full operation of the system. Shrink the primary volume by the required capacity, and then shrink the secondary volume.

Procedure

Complete the following steps to shrink a volume:

1. Validate that the volume is not mapped to any host objects. If the volume is mapped, data is displayed.
2. You can determine the exact capacity of the source or master volume. Issue the following command:

```
lsvdisk -bytes vdisk_name
```

3. Shrink the volume by the required amount. Enter the following command:

```
shrinkvdiskspace -size size_change -unit  
b | kb | mb | gb | tb | pb vdisk_name
```

where *size_change* indicates the size reduction for the volume in the specified units and *vdisk_name* is the volume that you are shrinking.

Migrating extents using the CLI

To improve performance, you can migrate extents using the command-line interface (CLI).

About this task

The system provides various data migration features. These features can be used to move the placement of data both *within* parent pools and *between* parent pools. These features can be used concurrently with I/O operations. You can use either of these methods to migrate data:

1. Migrating data (extents) from one MDisk to another (within the same parent pool). This method can be used to remove highly used MDisk.
2. Migrating volumes from one parent pool to another. This method can be used to remove highly used parent pools. For example, you can reduce the use of a pool of MDisk. Child pools that receive their capacity from parent pools, cannot have extents that are migrated to them.

Notes:

1. The source MDisk must not currently be the source MDisk for any other migrate extents operation.
2. The destination MDisk must not be the destination MDisk for any other migrate extents operation.

Migration commands fail if the target or source volume is offline, there is no quorum disk defined, or the defined quorum disks are unavailable. Correct the offline or quorum disk condition and reissue the command.

You can determine the use of particular MDisk by gathering input/output (I/O) statistics about nodes, MDisk, and volumes. After you collect this data, you can analyze it to determine which MDisk are used frequently. The procedure then takes you through querying and migrating extents to different locations in the same parent pool. This procedure only can be completed using the command-line interface.

If performance monitoring tools indicate that an MDisk in the pool is being overused, you can migrate data to other MDisk within the same parent pool.

Procedure

1. Determine the number of extents that are in use by each volume for the MDisk by issuing this CLI command:

```
lsmdiskextent mdiskname
```

This command returns the number of extents that each volume is using on the MDisk. Select some of these extents to migrate within the pool.

2. Determine the other MDisk that are in the same volume.
 - a. To determine the parent pool that the MDisk belongs to, issue this CLI command:

```
lsmdisk mdiskname | ID
```

- b. List the MDisk in the pool by issuing this CLI command:

```
lsmdisk -filtervalue mdisk_grp_name=mdiskgrpname
```

3. Select one of these MDisk as the target MDisk for the extents. You can determine how many free extents exist on an MDisk by issuing this CLI command:

```
lsfreeextents mdiskname
```

You can issue the **lsmdiskextent** *newmdiskname* command for each of the target MDisk to ensure that you are not just moving the over-utilization to another MDisk. Check that the 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 this CLI command to move them to another MDisk:

```
migrateexts -source mdiskname | ID -exts num_extents  
-target newmdiskname | ID -threads 4 -vdisk 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 this CLI command:

```
lsmigrate
```

Migrating volumes between pools using the CLI

You can migrate volumes between pools using the command-line interface (CLI).

About this task

You can determine the usage of particular MDisk by gathering input/output (I/O) statistics about nodes, MDisk, and volumes. After you collect this data, you can analyze it to determine which volumes or MDisk are hot. You can then migrate volumes from one storage pool to another.

Complete the following step to gather statistics about MDisk and volumes:

1. Use secure copy (**scp** command) to retrieve the dump files for analyzing. For example, issue the following command:

```
scp clusterip:/dumps/iostats/v_*
```

This command copies all the volume statistics files to the AIX host in the current directory.

2. Analyze the memory dumps to determine which volumes are hot. It might be helpful to also determine which MDisk are being used heavily as you can spread the data that they contain more evenly across all the MDisk in the storage pool by migrating the extents.

After you analyze the I/O statistics data, you can determine which volumes are hot. You also need to determine the storage pool that you want to move this volume to. Either create a new storage pool or determine an existing group that is not yet overly used. Check the I/O statistics files that you generated and then ensure that the MDisk or volumes in the target storage pool are used less than the MDisk or volumes in the source storage pool.

You can use data migration or volume mirroring to migrate data between storage pools. Data migration uses the command **migratevdisk**. Volume mirroring uses the commands **addvdiskcopy** and **rmvdiskcopy**.

Migrating data using `migratevdisk`

You can use the **`migratevdisk`** command to migrate data between two storage pools. When you issue the **`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 data migration function to move a volume between storage pools that have different extent sizes.
- Migration commands fail if the target or source volume is offline, there is no quorum disk defined, or the defined quorum disks are unavailable. Correct the offline or quorum disk condition and reissue the command.
- The system supports migrating volumes between child pools within the same parent pool or migrating a volume in a child pool to its parent pool. Migration of volumes fails if source and target child pools have different parent pools. However, you can use **`addvdiskcopy`** and **`rmvdiskcopy`** commands to migrate volumes between child pools in different parent pools.

When you use data migration, it is possible for the free destination extents to be consumed by another process; for example, if a new volume is created in the destination parent pool 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 more MDiskS to the target parent pool, which provides more 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 volumes that are already created from the parent pool to another group. This action frees up extents in the group and allows the original migrations to be restarted.

Complete the following steps to use the **`migratevdisk`** command to migrate volumes between storage pools:

1. After you determine the volume that you want to migrate and the new storage pool that you want to migrate it to, issue the following CLI command:

```
migratevdisk -vdisk vdisk_name
             -mdiskgrp
             mdisk_group_name -threads 4
```

2. You can check the progress of the migration by issuing the following CLI command:

```
lsmigrate
```

Migrating data using volume mirroring

When you use data migration, the volume goes offline if either pool fails. Volume mirroring can be used to minimize the impact to the volume because the volume goes offline only if the source pool fails. You can migrate volumes between child pools or from a child pool to a parent pool using the **`addvdiskcopy`** and **`rmvdiskcopy`** commands instead of using the **`migratevdisk`** command. Complete the following steps to use volume mirroring to migrate volumes between pools:

1. After you determine the volume that you want to migrate and the new pool that you want to migrate it to, enter the following command:

```
addvdiskcopy -mdiskgrp mdisk_group_name vdisk_name -autodelete yes
```

where *mdisk_group_name* is the name of the new storage pool and *vdisk_name* is the name of the volume that is being copied. Specify *yes* for the **`-autodelete`** to automatically delete the original copy of the volume after the copies are synchronized.

2. The copy ID of the new copy is returned. The copies now synchronize such that the data is stored in both storage pools. You can check the progress of the synchronization by issuing the following command:

```
lsvdisksyncprogress
```

Moving a volume between I/O groups using the CLI

To move volumes between I/O groups non-disruptively, ensure that hosts are mapped to the volume, support non-disruptive volume move. The cached data that is held within the system must first be written to the system disk before the allocation of the volume can be changed.

About this task

Modifying the I/O group that services the volume can be done concurrently with I/O operations if the host supports non-disruptive volume move. 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 volume is accessed has changed. This can be done in the situation where one pair of nodes has become over used.

If there are any host mappings for the volume, the hosts must be members of the target I/O group or the migration fails.

Make sure you create paths to I/O groups on the host system. After the system has successfully added the new I/O group to the volume's access set and you have moved selected volumes to another I/O group, detect the new paths to the volumes on the host. The commands and actions on the host vary depending on the type of host and the connection method used. These steps must be completed on all hosts to which the selected volumes are currently mapped.

You can also use the management GUI to move volumes between I/O groups non-disruptively. In the management GUI, select **Volumes > Volumes**. On the Volumes panel, select the volume that you want to move and select **Actions > Move to Another I/O Group**. The wizard guides you through all the steps that are necessary for moving a volume to another I/O group, including any changes to hosts that are required. Click **Need Help** on the associated management GUI panels for details.

Note: If the selected volume is performing quick initialization, this wizard is unavailable until quick initialization is complete.

To move a volume between I/O groups using the CLI, complete the following steps:

Procedure

1. Issue the following command: **addvdiskaccess -iogrp iogrp id/name volume id/name**
2. Issue the following command: **movevdisk -iogrp destination iogrp -node new preferred node volume id/name** The system disables moving a volume if the selected volume is currently performing quick initialization. After the quick initialization completes, you can move the volume to another I/O group.
3. Issue the appropriate commands on the hosts mapped to the volume to detect the new paths to the volume in the destination I/O group.
4. After you confirm the new paths are online, remove access from the old I/O group: **rmvdiskaccess -iogrp iogrp id/name volume id/name**
5. Issue the appropriate commands on the hosts mapped to the volume to remove the paths to the old I/O group.

Creating an image-mode 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 volumes.

About this task

Make sure you are aware of the following information before you create image-mode 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 clustered system by adding these disks one at a time. For example, map a single LUN from your RAID storage system to the clustered system 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 a parent pool. If you do, the data is lost. When you use the command to create an image-mode volume from an unmanaged-mode disk, select the parent pool where it should be added. Ensure that the pool that is selected is not a child pool. Child pools are created from existing pools, called parent pools, and get capacity from the parent pool, not MDisks.

Complete the following steps to create an image-mode volume:

Procedure

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 storage pools. Ensure that the pool is not a child pool.
3. Map a single array or logical unit from your RAID storage system to the clustered system. You can do this through a switch zoning or a RAID storage system based on your host mappings. The array or logical unit appears as an unmanaged-mode MDisk to the system.
4. Issue the **lsmdisk** command to list the unmanaged-mode MDisks.
If the new unmanaged-mode MDisk is not listed, you can complete a fabric-level discovery. Issue the **detectmdisk** command to scan the Fibre Channel network for the unmanaged-mode MDisks.

Note: The **detectmdisk** command also rebalances MDisk access across the available storage system device ports.

5. Convert the unmanaged-mode MDisk to an image-mode volume.

Note: If the volume that you are converting maps to a flash drive, the data that is stored on the volume is not protected against Flash drive failures or node failures. To avoid data loss, add a volume copy that maps to an Flash drive on another node.

Issue the **mkvdisk** command to create an image-mode volume object.

6. Map the new volume to the hosts that were previously using the data that the MDisk now contains. You can use the **mkvdiskhostmap** command to create a new mapping between a volume and a host. This makes the image-mode volume accessible for I/O operations to the host.

Results

After the volume is mapped to a host object, the volume is detected as a disk drive with which the host can complete I/O operations.

What to do next

If you want to virtualize the storage on an image-mode volume, you can transform it into a striped volume. Migrate the data on the image-mode volume to managed-mode disks in another storage pool. Issue the **migratevdisk** command to migrate an entire image-mode volume from one storage pool to

another storage pool. Ensure that the storage pool that you migrate the image-mode volume to is not a child pool.

Migrating data to an image mode volume using the CLI

You can use the command-line interface (CLI) to migrate data to an image mode volume.

About this task

Use the **migratetoimage** CLI command to migrate the data from an existing volume onto a different managed disk (MDisk).

When the **migratetoimage** CLI command is issued, it migrates the data of the user specified source volume onto the specified target MDisk. When the command completes, the volume is classified as an image mode volume.

Note: Migration commands fail for the following reasons:

- The target or source volume is offline.
- A quorum disk is not defined.
- Defined quorum disks are unavailable.

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. Using this command results in the inclusion of the MDisk into the user specified storage pool.

Enter the following CLI command to migrate data to an image mode volume:

```
migratetoimage -vdisk source_vdisk_name -mdisk unmanaged_target_mdisk_name -mdiskgrp managed_disk_group_name
```

where *source_vdisk_name* is the name of the image mode volume, *unmanaged_target_mdisk_name* is the name of the new MDisk, and *managed_disk_group_name* is the name of the new storage pool. For example, the following command migrates data from the `vdisk0` image mode volume to the `mdisk5` target in the `mdgrp2` storage pool:

```
migratetoimage -vdisk vdisk0 -mdisk mdisk5 -mdiskgrp mdgrp2
```

Deleting a node from a system by using the CLI

You can use the command-line interface (CLI) to remove a node from a system.

Before you begin

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 system 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 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 destroys the system. Before you delete the last node in the system, ensure that you want to destroy the system.
- 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. These 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 volumes are assigned to the I/O group, you cannot delete the node from the system 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.

About this task

Complete these steps to delete a node:

Procedure

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 this CLI command to request a filtered view of the volumes:

```
lsvdisk -filtervalue IO_group_name=name
```

Where *name* is the name of the I/O group.

- b. Issue this CLI command to list the hosts that this volume is mapped to:

```
lsvdiskhostmap vdiskname/identification
```

Where *vdiskname/identification* is the name or identification 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 clustered system, 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 clustered system. This process might take a short time, typically less than a minute. The system IP address remains unchanged, but any SSH client that is 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 system. The system causes the node to remove itself from the system and the node becomes a candidate to add to this system or another system.
 - c. If you are adding this node into the system, 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 system 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 host operating system.

4. Issue this CLI command to delete a node from the clustered system:

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:

- a. Run the **lsdependentvdisks** command.
- b. For each node-dependent volume that is returned, run the **lsvdisk** command.
- c. Ensure that each volume returns in-sync status.

```
rmnode node_name_or_identification
```

Where *node_name_or_identification* is the name or identification of the node.

Note: Before removing a node, the command checks for any node-dependent volumes that would go offline. If the node that you selected to delete contains a flash drive that has dependent volumes, volumes that use the flash drives go offline and become unavailable if the node is deleted. To maintain access to volume data, mirror these volumes before removing the node. To continue removing the node without mirroring the volumes, specify the **force** parameter.

Completing the system maintenance procedure by using the CLI

You can use the command-line interface (CLI) to complete the system maintenance procedure.

About this task

Use the following steps to complete the system maintenance procedure:

Procedure

1. Issue the **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 order that is defined within the code, the highest priority of unfixed errors is returned.
2. Issue the **dumperrlog** command to dump the contents of the error log to a text file.
3. Locate and fix the error.

4. Issue the **clearerrlog** command to clear all entries from the error log, including status events and any unfixed errors. Only issue this command when you rebuild the system or you fix a major problem that 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 **cherrstate** command to toggle the state of an error between unfixed and fixed.

Modifying system IP addresses using the CLI

Use the command-line interface (CLI) to change the IP addresses that are associated with a system.

About this task

Attention: When you specify a new IP address for a system, the existing communication with the system is broken. You must reconnect to the system with the new IP address. Additionally, the address for a system IP cannot be the same address used for the service IP. Using the same IP address causes an error.

Procedure

To change the system IP address, complete the following steps:

1. Issue the **lsystemip** command to list IP addresses that are used by the system.
2. Record the IP addresses for future reference.
3. To change an Internet Protocol Version 4 (IPv4) system IP address, issue this command:
`chsystemip -clusterip cluster_ip_address -port cluster_port`
where *cluster_ip_address* is the new IP address for the system and *cluster_port* specifies the port (1 or 2) where changes apply.
4. To change an IPv4 system IP address to an IPv6 system IP address, issue this command:
`chsystemip -clusterip_6 cluster_ip_address -port cluster_port`
where *cluster_ip_address* is the new Internet Protocol Version 6 (IPv6) address for the system and *cluster_port* specifies the port (1 or 2) where changes apply.
5. To change an IPv4 default gateway IP address, issue this command:
`chsystemip -gw cluster_gateway_address -port cluster_port`
where *cluster_gateway_address* is the new gateway address for the system and *cluster_port* specifies the port (1 or 2) where changes apply.
6. To change an IPv6 default gateway address, issue this command:
`chsystemip -gw_6 cluster_gateway_address -port cluster_port`
where *cluster_gateway_address* is the new gateway address for the system and *cluster_port* specifies the port (1 or 2) where changes apply.
7. Issue this command to change an IPv4 system subnet mask
`chsystemip -mask cluster_subnet_mask -port cluster_port`
where *cluster_subnet_mask* is the new subnet mask for the system and *cluster_port* specifies the port (1 or 2) where changes apply.
8. For IPv6 addresses, you can issue this command to set the prefix for the system:
`chsystemip -prefix_6 -port cluster_port`
where *cluster_port* specifies the port (1 or 2) where changes apply.
9. Optionally, to delete all of the IPv4 addresses in the system after you change all addresses to IPv6, issue this command:
`chsystem -noip`

10. Optionally, to delete all of the IPv6 addresses in the system after you change all addresses to IPv4, issue this command:

```
chsystem -noip_6
```

11. Display the IP routing table (optional) by using the CLI command `lsroute` as follows:

```
lsroute
```

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.

12. Issue the **ping** command (optional) to diagnose IP configuration problems. An example of the command is as follows:

```
ping -srcip4 source_ipv4_address destination_ipv4_address -srcip6 source_ipv6_address destination_ipv6_address
```

Changing the system gateway address by using the CLI

You can use the command-line interface (CLI) to change the gateway address for a system.

Procedure

To change the system gateway address, complete the following steps.

1. Issue the **lssystemip** command to list the current gateway address of the system.
2. Record the current gateway address for future reference.
3. Issue the following command to change an IPv4 system gateway address:

```
chsystemip -gw cluster_gateway_address -port cluster_port
```

where *cluster_gateway_address* is the new gateway address for the system. The **port** parameter specifies which port (1 or 2) to apply changes to.

4. Issue the following command to change an IPv6 system gateway address:

```
chsystemip -gw_6 cluster_gateway_address -port cluster_port
```

where *cluster_gateway_address* is the new gateway address for the system. The **port** parameter specifies which port (1 or 2) to apply changes to.

Changing the relationship bandwidth for a system by using the CLI

You can use the command-line interface (CLI) to change the relationship bandwidth for a system.

About this task

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 system partnership. The default value for the relationship bandwidth limit is 25 megabytes per second (MBps), but you can change this value by following these steps:

Procedure

1. Issue the **lssystem** command to list the current relationship bandwidth limit of the system. For example:

```
lssystem system_id_or_system_name
```

Where *system_id_or_system_name* is the ID or name of the system.

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 system, issue the following command:

```
chsystem -relationshipbandwidthlimit  
system_relationship_bandwidth_limit
```

Where *system_relationship_bandwidth_limit* is the new limit for the system. Issue the command on both systems in a relationship.

Configuring the system for iSCSI hosts

You need to complete several tasks to configure the system to work with iSCSI-attached hosts. The tasks include general tasks on the host system before you configure a system.

Before you begin

Before you complete any iSCSI-configuration tasks on the system, it is important that you complete all the iSCSI-related configuration on the host machine. Because the system supports various host machines, consult the documentation for specific instructions and requirements for a particular host. For a list of supported hosts, see this website:

www.ibm.com/support

About this task

To configure a system 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 system. Hosts use iSCSI names to connect to the node. Each node, for example, has a unique IQN, and the system name and node name are used as part of that IQN. Each node, for example, has a unique IQN, and the system 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 process I/O.

Procedure

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:

```
cfgportip -node node_name | node_id -ip ipv4addr  
-gw ipv4gw -mask subnet_mask -failover -vlan vlan_id 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 **lspportip** command.

The optional **-failover** parameter specifies that the IP is a failover IP and 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 the partner node.

The optional **-vlan** parameter sets the virtual local area network (VLAN) ID for an IPv4 address that is configured for iSCSI host attachment.

For details regarding the **-vlan** parameter, see the information about configuring VLAN for iSCSI using the CLI.

2. To configure a new port IP address to a specified Ethernet port of a node with an IPv6 address, enter the following CLI command:

```
cfgportip -node node_name | node_id -ip_6 ipv6addr  
-gw_6 ipv6gw -prefix_6 prefix -failover -vlan_6 vlan_id 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 Ethernet port, *ipv6gw* is the IPv6 gateway IP address, *subnet_mask* is the IPv6 subnet mask, and *port_id* specifies the Ethernet port ID (1 or 2). To view a list of ports, use the **lspportip** command.

The optional **-failover** parameter specifies that the IP is a failover IP 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 the node that is specified.

The optional **-vlan_6** parameter sets the virtual local area network (VLAN) ID for an IPv6 address that is configured for iSCSI host attachment.

For details regarding the **-vlan** parameter, see the information about configuring VLAN for iSCSI using the CLI.

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:

```
rmportip -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 **lspportip** 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:

```
rmportip -ip_6 -failover  
-node node_name | node_id port_id
```

where **-ip_6** indicates that this command removes 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 **lspportip** command. The optional **-failover** parameter indicates that the specified data is failover data.

What to do next

After you configure your IP addresses, you can optionally create or configure the following iSCSI items:

Configuring or modifying an iSCSI alias by 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.

About this task

To configure or modify an iSCSI alias, follow these steps:

Procedure

1. To configure a new port IP address to a specified Ethernet port of a node, enter the following CLI command:

```
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 clustered system. If this parameter is used when there is a partner node, the name or alias of that node changes

```
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.

What to do next

After you create iSCSI aliases, you can optionally configure the address for the Internet Storage Name Service (iSNS) server for the system.

Configuring the iSNS server address using the CLI

If you are using iSCSI-attached hosts with the SAN Volume Controller clustered system, you can use the command-line interface (CLI) to optionally configure the address for the Internet Storage Name Service (iSNS) server for the system. Host systems use the iSNS server to manage iSCSI targets and for iSCSI discovery.

Procedure

1. To specify an IPv4 address for the iSCSI storage name service (SNS), enter the following CLI command:

```
chsystem -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:

```
chsystem -isnsip_6 ipv6_sns_server_address
```

where *ipv6_sns_server_address* is the IP address of the iSCSI storage name service in IPv6 format.

Configuring system iSCSI authentication using the CLI

You can use the command-line interface (CLI) to configure the SAN Volume Controller clustered system to authenticate with iSCSI-attached hosts using the Challenge-Handshake Authentication Protocol (CHAP). After the CHAP is set for the system, all attached hosts must be configured to authenticate. When troubleshooting a problem, you can delay your configuration of the CHAP authentication until after you configure the first one or two hosts and test their connectivity.

About this task

To configure authentication between the SAN Volume Controller system and the iSCSI-attached hosts, follow these steps:

Procedure

1. To configure CHAP authentication for an iSCSI host, enter the following CLI command:

```
chhost -chapsecret chap_secret host_name
```

where *chap_secret* is the CHAP secret to be used to authenticate the system via iSCSI and *host_name* is the name of the iSCSI host. The *chap_secret* value must be 12 characters.

2. To set the authentication method for the iSCSI communications of the system, enter the following CLI command:

```
chsystem -iscsiauthmethod chap -chapsecret chap_secret
```

where *chap* specifies that CHAP is the authentication method and *chap_secret* is the CHAP secret to be used. The specified CHAP secret cannot begin or end with a space.

3. To clear all CHAP secrets for iSCSI authentication that have been previously set, enter the following CLI command:

```
chsystem -nochapsecret
```

The **nochapsecret** parameter is not allowed if the **chapsecret** parameter is specified.

4. Run the **lsiscsiauth** command to display the Challenge Handshake Authentication Protocol (CHAP) secret that has been configured.

What to do next

After you configure the CHAP secret for the SAN Volume Controller system, ensure that the system 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 system.

Configuring remote authentication service by using the CLI

You can use the command-line interface (CLI) to configure the system to use remote authentication.

About this task

If a user is configured on the system as a local user, only local credentials are used. Otherwise, users who are entering their passwords are authenticated against the remote service when they use the management GUI or the command-line interface (CLI). Their roles are determined according to group memberships defined on the remote service. If a user is configured on the system as a remote user with an SSH key, the user can additionally access the command-line interface by using this Secure Shell (SSH) key. Group memberships continue to be determined from the remote service.

Configuring remote authentication service with Lightweight Directory Access Protocol (LDAP) using the CLI

You can use the command-line interface (CLI) to configure the system to authenticate users against servers that implement the Lightweight Directory Access Protocol (LDAP), including IBM Security Services and Active Directory (AD).

About this task

- Users on provisioned LDAP servers with IBMRBS permissions of Supervisor Access or Supervisor Role can log in to the system as Administrator, but cannot run the **satask** command.
- All authentication commands and settings are disabled.
 - Automatically provisioned settings are not visible to the user and are not displayed by the **lssystem** or **lsldapserver** commands.
 - The **chauthservice -refresh** command is enabled.

All options on the system GUI LDAP page are disabled.

Tip: A superuser cannot be authenticated if the superuser is using a remote Lightweight Directory Access Protocol (LDAP server). However, other users can authenticate in this manner.

Procedure

To enable user authentication with LDAP, follow these steps:

1. Configure LDAP by issuing the **chldap** command.

This command provides default settings for both Tivoli Directory Server and AD. To configure authentication with Tivoli Directory Server schema defaults and Transport Layer Security (TLS), for example, issue the following command:

```
chldap -type itds -security tls
```

LDAP configuration can be inspected with the **lsldap** command.

Note: TLS is recommended because transmitted passwords are encrypted.

2. Specify the **mkldapserver** command to define up to six LDAP servers to use for authentication.

Multiple servers can be configured to provide access to different sets of users or for redundancy. All servers must share the settings that are configured with **chldap**. To configure an LDAP server with a Secure Socket Layer (SSL) certificate and users in the `cn=users,dc=company,dc=com` subtree, for example, issue:

```
mkldapserver -ip 9.71.45.108 -basedn cn=users,dc=company,dc=com -sslcert /tmp/sslcert.pem
```

You can also configure which servers are preferred to authenticate users.

Specify **lsldapserver** for LDAP server configuration information. Specify **chldapserver** and **rmldapserver** to change the configured LDAP servers.

3. Configure user groups on the system by matching those user groups that are used by the authentication service.

For each group of interest that is known to the authentication service, a system user group must be created with the same name and with the remote setting enabled. If members of a group that is called `sysadmins`, for example, require the system administrator (admin) role, issue the following command:

```
mkusergrp -name sysadmins -remote -role Administrator
```

If none of the user groups match a system user group, the user cannot access the system.

4. Verify your LDAP configuration by using the **testldapserver** command.

To test the connection to the LDAP servers, issue the command without any options. A user name can be supplied with or without a password to test for configuration errors. To process a full authentication attempt against each server, issue the following commands:

```
testldapserver -username username -password password
```

5. Issue the following command to enable LDAP authentication:

```
chauthservice -type ldap -enable yes
```

6. Configure users who do not require Secure Shell (SSH) key access.

Delete system users who must use the remote authentication service and do not require SSH key access.

Remember: A superuser cannot be deleted or use the remote authentication service.

7. Configure users who require SSH key access.

All system users who use the remote authentication service and require SSH key access must have remote settings that are enabled and a valid SSH key that is configured on the system.

Changing user groups

You can use the command-line interface (CLI) to change user groups. User groups organize users of a clustered system by role.

About this task

Roles apply to both local and remote users on the system and are based on the user group to which the user belongs. A local user can belong only 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 change a user group in the management GUI, select **Access > Users**. Select a user group and select **Properties** from the **Actions** menu.

Procedure

1. Issue the **chusergrp** CLI command to change attributes of an existing user group. For example:

```
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.

2. Issue the **lsusergrp** CLI command to display the user groups that were created on the system. For example:

```
lsusergrp 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 system are displayed.

Changing users

You can use the command-line interface (CLI) or the management GUI to change users on the system.

Before you begin

System users must provide either a password, a Secure Shell (SSH) key, or both. Local users are authenticated through the authentication methods that are on the system.

You can create two categories of users that access the clustered system (system). These user types are based on how they authenticate to the system:

- Some users must provide an SSH password (or if not possible an SSH key).
- If a 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 password and SSH key can be used.
- Users must be part of a user group that is defined on the system.

Remote users should also configure local credentials if they need to access the system when the remote service is down. Remote users have their groups that are defined by the remote authentication service.

To change a user in the management GUI, select **Access > Users**. Right-click the user and select **Modify** from the **Actions** menu.

About this task

To change a user in the CLI, follow these steps:

Procedure

1. Issue the **chuser** CLI command to change the attributes of an existing user. For example:

```
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.

2. Issue the **chcurrentuser** CLI command to change the attributes of the current user. For example:

```
chcurrentuser -nokey
```

where the **nokey** parameter specifies that the SSH key of the user is to be deleted.

3. Issue the **lscurrentuser** CLI command to display the name and role of the logged-in user. For example:

1currentuser

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).

About this task

The notification settings apply to the entire system. You can specify the types of events that cause the system to send a notification. The system 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. If no community string is specified, the default string of **public** is used. 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 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:

Procedure

1. To create a new SNMP server to receive notifications, use the **mksnmpserver** CLI command. For example, enter one of the following commands:

```
mksnmpserver -ip 9.11.255.634
```

where *9.11.255.634* is the IP address for this server.

```
mksnmpserver -ip 9.11.255.634 -port remoteportnumber
```

where *9.11.255.634* is the IP address for this server and *remotepor~~tn~~umber* is the port number for the remote SNMP server.

2. To change the settings of an existing SNMP server, enter the **chsnmpserver** command. For example:

```
chsnmpserver -name server_name snmp_server_name_or_id
```

where *server_name* is the new name of the server and *snmp_server_name* is the name or ID of the server to be modified.

3. To remove an existing SNMP server from the system, enter the **rmsnmpserver** command. For example:

```
rmsnmpserver snmp_server_name
```

where *snmp_server_name* is either the name 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 system, enter the **lssnmpserver** command. For example, to display a concise view, enter the following command:

```
lssnmpserver -delim :
```

To display a detailed view of an SNMP server, enter the following command:

Setting up syslog notifications using the CLI

You can set up syslog event notifications using the command-line interface (CLI).

About this task

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 specify up to a maximum of six syslog servers. You can use the management GUI or the command-line interface to configure and modify your syslog settings.

The syslog event notification settings apply to the entire system. You can specify the types of events that cause the system 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.

To configure and work with notification settings, use the following commands:

Procedure

1. Issue the **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:

```
mksyslogserver syslog_server_name -ip 9.11.255.123
```

where *syslog_server_name* is the name given to the Syslog server definition and *9.11.255.123* is the external Internet Protocol (IP) address of the syslog server.
2. To modify a syslog notification, issue the **chsyslogserver** command. For example:

```
chsyslogserver syslog_server_name -ip 9.11.255.123
```

where *syslog_server_name* is the name given to the Syslog server definition and *9.11.255.123* is the external IP address of the syslog server.
3. To delete a syslog notification, issue the **rmsyslogserver** command. For example:

```
rmsyslogserver syslog_server_name -force
```
4. To display either a concise list or a detailed view of syslog servers that are configured on the system, issue the **lssyslogserver** command. For example, to display a concise view, enter the following command:

```
lssyslogserver -delim :
```

To display a detailed view of a syslog server, enter the following command:

```
lssyslogserver syslog_server_name
```

Setting up email 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.

Before you begin

You can configure Call Home using the CLI.

About this task

To set up, manage, and activate email event, inventory, and Call Home notifications, complete the following steps:

Procedure

1. Enable your system to use the email notification function. To do this, issue the **mkemailserver** CLI command. Up to six SMTP email servers can be configured to provide redundant access to the external email network.

This example creates an email server object. It specifies the name, IP address, and port number of the SMTP email server. After you issue the command, you see a message that indicates that the email server was successfully created.

```
mkemailserver -ip ip_address -port port_number
```

where *ip_address* specifies the IP address of a remote email server and *port_number* specifies the port number for the email server.

2. Add recipients of email event and inventory notifications to the email event notification facility. To do this, issue the **mkemailuser** CLI command.

The following example adds email recipient *manager2008* and designates that this recipient is to receive email error-type event notifications.

```
mkemailuser -address manager2008@ibm.com  
-error on -usertype local
```

Important: Always select the *local* usertype unless otherwise instructed by IBM Support. The *support* usertype is normally only used with the Call Home feature.

Remember: To control the frequency of email notifications, issue the following command: **chsystem** -inventoryemail

3. Set the contact information that is used by the email event notification facility. To do this, issue the **chemail** CLI command. If you are starting the email event notification facility, the **-reply** parameter must be set. I

The following example sets the contact information for the email recipient *manager2008*.

```
chemail -reply manager2008@ibm.com -contact manager2008  
-primary 0441234567 -location 'room 256 floor 1 IBM'
```

4. Optionally, generate a report that lists email event notification settings for all email recipients, or change or delete email recipients.
 - To generate a report that lists the email event notification settings for all email recipients, an individual email recipient, or a specified type of email recipient (local or support), issue the **lsemailuser** CLI command.
 - To change the settings that are defined for a recipient, issue the **chemailuser** CLI command. You must specify the user ID or name of the email recipient for whom you are modifying settings.
 - To remove a previously defined email recipient, issue the **rmemailuser** CLI command. You must specify the user ID or name of the email recipient that you want to remove.
5. Activate the email and inventory notification function. To do this, issue the **startemail** CLI command. There are no parameters for this command.

Note: Inventory information is automatically reported to service personnel when you activate error reporting.

6. Optionally, test the email notification function to ensure that it is operating correctly and send an inventory email notification. SAN Volume Controller uses the notifications settings to call home if errors occur.
 - To send a test email notification to one or more recipients, issue the **testemail** CLI command. You must either specify **all** or the user ID or user name of an email recipient that you want to send a test email to.
 - To send an inventory email notification to all recipients that are enabled to receive inventory email notifications, issue the **sendinventoryemail** CLI command. There are no parameters for this command.
 - Use the **stopemail** command to stop the email and inventory notification function. There are no parameters for this command.

Setting up email servers using the CLI

You can set up email server objects by using the command-line interface (CLI).

About this task

You can specify a server object that describes a remote Simple Mail Transfer Protocol (SMTP) email server to receive event notifications from the clustered system. You can specify up to six servers to receive notifications. To configure and work with email servers, use the following commands:

Procedure

1. Issue the **mkemailserver** CLI command to create an email server object that describes a remote Simple Mail Transfer Protocol (SMTP) email server. For example, issue the following CLI command to set up an email server:

```
mkemailserver -ip ip_address
```

where *ip_address* is the IP address of a remote email server. This address must be a valid IPv4 or IPv6 address.
2. To change the parameters of an existing email server object, issue the **chemailserver** command. For example:

```
chemailserver -ip ip_address email_server_name_or_id
```

where *ip_address* is the IP address of the email 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 email server object, issue the **rmemailserver** command. For example:

```
rmemailserver email_server_name_or_id
```
4. To display either a concise list or a detailed view of email servers that are configured on the system, issue the **lserver** command. For example, to display a concise view, enter the following command:

```
lserver -delim :
```

To display a detailed view of an email server, enter the following command:

```
lserver email_server_name_or_id
```

Changing user passwords using the CLI

You can use the command-line interface (CLI) to change user passwords.

About this task

Passwords control access to these applications:

- System management GUI

- Service assistant GUI
- CLI

Follow these steps to change the password for a user:

Procedure

Enter the following command to change the password:

```
chuser -password cleartextpassword janedoe
```

Where *password* is the new password that you want to use for the user janedoe.

Changing the locale setting using the CLI

You can use the command-line interface (CLI) to specify the locale for a system. The language that you select as your locale setting is used to display command results and error messages in the CLI.

About this task

The following locales are available:

- 0 US English (default)
- 3 Japanese

Procedure

Issue the **setlocale** CLI command with the ID for the locale.

Example

For example, issue the following CLI command to change the locale setting from US English to Japanese:

```
setlocale -locale 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.

About this task

Perform the following steps to view the feature log:

Procedure

1. Issue the **lsdumps** 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 **lsdumps** 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).

About this task

Perform the following step to analyze the error log:

Procedure

Issue the following CLI command to list error log entries by file type: **lseventlog**

Results

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 system by using the CLI

You can use the command-line interface (CLI) to shut down a system.

Procedure

To power off your system, complete the following steps.

1. Determine which hosts have access to volumes on this system by running the **lshostvdiskmap** command.
2. Stop input/output (I/O) to the system from each host that is listed in step 1.

Note: Failure to stop host I/O can result in failed I/O operations being reported to your host operating systems.

3. Shut down the system by using this command:
`stopssystem`
4. Wait for the power light-emitting diodes (LEDs) on all node canisters in all control enclosures to flash at 1 Hz, indicating that the shutdown operation has completed.




Figure 1. Location of the power LED on a node canister

- 1** Power
- 2** Status
- 3** Fault

5. Disconnect the power cords from both power supplies in each control enclosure.
6. Disconnect the power cords from both power supplies in each expansion enclosure.

Updating the software automatically using the CLI

You can use the command-line interface (CLI) to install software updates.

Before you begin

Follow these steps to update automatically by using the CLI to version 7.5 or later from version 7.4 or later. If you are updating from a release previous to version 7.4.0, follow the instructions in that previous release. You are required, however, to confirm the update, which is not included in those steps. After you follow the instructions for your release, return here and continue with step 8 on page 86 in the procedure.

To update from version 5.1.x or earlier, see the relevant IBM Knowledge Center or publications that are available at this website: www.ibm.com/support

Important: Before you start an update, you must check for offline or degraded volumes. An offline volume can cause write data that was modified to be pinned in the system cache. This action prevents volume failover and causes a loss of input/output (I/O) access during the update. If the `fast_write_state` is empty, a volume can be offline and not cause errors during the update.

If you encounter a memory DIMM failure to any node during the update process, stop immediately and follow directions in Updating the system.

About this task

To update the system, follow these steps.

Procedure

1. You must download, install, and run the latest version of the test utility to verify that there are no issues with the current system.

Important: After you have installed and run the test utility using either the GUI or the CLI, you must return to step 2 on this Knowledge Center page.

You can download and install the most current version of this tool at the following website. The link provides instructions for using either the GUI or the CLI to install and run the test utility.

<http://www.ibm.com/support/docview.wss?uid=ssg1S4000585>

2. Download the latest code from the www.ibm.com/support site.
 - If you want to write the code to a CD, you must download the CD image.
 - If you do not want to write the code to a CD, you must download the installation image.
3. Use PuTTY scp (pscp) to copy the update files to the node.
4. Ensure that the update file was successfully copied.

Before you begin the update, you must be aware of the following situations:

- The installation process fails under the following conditions:
 - If the code that is installed on the remote system is not compatible with the new code or if an intersystem communication error does not allow the system to check that the code is compatible.
 - If any node in the system has a hardware type that is not supported by the new code.
 - If the system determines that one or more volumes in the system would be taken offline by rebooting the nodes as part of the update process. You can find details about which volumes would be affected by using the **lsdependentvdisks** command. If you are prepared to lose access to data during the update, you can use the force flag to override this restriction.
- The update is distributed to all the nodes in the system by using internal connections between the nodes.
- Nodes are updated one at a time.

- Nodes run the new code concurrently with normal system 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 thirty-minute delay between node updates. The delay allows time for the host multipathing software to rediscover paths to the nodes that are updated. There is no loss of access when another node in the I/O group is updated.
- The update is not committed until all nodes in the system are successfully updated to the new code level. If all nodes are successfully restarted with the new code level, the new level is committed. When the new level is committed, the system vital product data (VPD) is updated to reflect the new code level.
- Wait until all member nodes are updated and the update is committed before you invoke the new functions of the updated code.
- Because the update process takes some time, the installation command completes as soon as the code level is verified by the system. To determine when the update is completed, you must either display the code level in the system VPD or look for the **Software update complete** event in the error/event log. If any node fails to restart with the new code level or fails at any other time during the process, the code level is backed off.
- During an update, the version number of each node is updated when the code is installed and the node is restarted. The system code version number is updated when the new code level is committed.
- When the update starts, an entry is made in the error or event log and another entry is made when the update completes or fails.

5. Issue this CLI command to start the update process:

```
applysoftware -file software_update_file
```

Where *software_update_file* is the name of the code update file in the directory you copied the file to in step 3 on page 85. If the system identifies any volumes that would go offline as a result of rebooting the nodes as part of the system update, the code update does not start. An optional **force** parameter can be used to indicate that the update continues regardless of the problem identified. If you use the **force** parameter, you are prompted to confirm that you want to continue. The behavior of the **force** parameter changed, and it is no longer required when you apply an update to a system with errors in the event log.

6. If you are updating from a release before version 7.4.0, issue the following CLI command to check the status of the code update process:

```
svcinfolsoftwareupgradestatus
```

This command displays inactive when the update is complete.

Note: If a status of stalled_non_redundant is displayed, proceeding with the remaining set of node updates might result in offline volumes. Contact a service representative to complete the update.

7. If you are updating from version 7.4.0 or later, issue the following CLI command to check the status of the code update process:

```
lsupdate
```

This command displays success when the update is complete.

Note: If a status of stalled_non_redundant is displayed, proceeding with the remaining set of node updates might result in offline volumes. Contact a service representative to complete the update.

8. If you updated from a release before version 7.4.0, you receive the status message `system_completion_required`. To complete the update process, issue the command **applysoftware -complete**. After that command is run, you can run **lsupdate** to see the progress of the update completion.

9. To verify that the update successfully completed, issue the **lsnodevdpd** CLI command for each node that is in the system. The code version field displays the new code level.

Results

When a new code level is applied, it is automatically installed on all the nodes that are in the system.

Note: An automatic system update can take up to 30 minutes per node.


Chapter 3. Array commands

Array commands capture information that can assist you with managing arrays.

charray

Use the **charray** command to change array attributes.

Syntax



Parameters

-name new_name_arg

(Optional) Specifies the new name to apply to the array MDisk.

-sparegoal 1-100

(Optional) Sets the number of spares to protect the array members with. The value can be a number between 1 and 100.

Note: This parameter is not applicable for distributed arrays.

-balanced

(Optional) Forces the array to balance and configure the spare goals of the present drives.

Specify **-balanced** and the system examines the membership's chain balance for mirrored arrays. If each mirrored member is on a different chain than its partner member, the array continues balancing the member chains. If each mirrored member is not on a different chain than its partner member, the array stops balancing the member chains.

Note:

- If **-balanced** is specified and the goal of the associated array MDisk changes, the tier of the array MDisk is updated to match the new goal.
- This parameter is not applicable for distributed arrays.

-slowwritepriority latency | redundancy

(Optional) Controls array ability to complete write operations that take too long, even if it temporarily compromises redundancy.

The value can be either **latency** or **redundancy**:

- **latency** implies the feature is enabled for normal I/O operations
- **redundancy** implies the feature is not enabled for normal I/O operations

The default value is **latency** mode for existing arrays, unless the array is RAID-0 (in which case **redundancy** mode is required).

Important: Do not change the mode of a RAID-0 array.

Important: An array can cause member drives to become unsynchronized (to preserve response time) if the value is latency. If the value is redundancy, the array cannot cause member drives to become unsynchronized (to preserve time) and I/O performance is impacted.

-rebuildareagoal 0 | 1 | 2 | 3 | 4

(Optional) Specifies the rebuild areas threshold. The array logs an error when the available rebuild areas drop below this specified threshold. The values are 0, 1, 2, 3, or 4. (If you specify 0, an error is not logged if the system runs out of rebuild areas.)

Note: This parameter is only applicable for distributed arrays.

mdisk_id | *mdisk_name*

(Required) Identifies (by ID or user-defined name) which array the MDisk command applies to.

Description

This command changes an array's attributes.

An invocation example to change the name of an array

```
charray -name raid6 mdisk0 0
```

The resulting output:

No feedback

An invocation example to set the number of spares threshold to 2

```
charray -sparegoal 2 mdisk52
```

The resulting output:

No feedback

An invocation example to balance the array

```
charray -balanced 3
```

The resulting output:

No feedback

An invocation example for changing the rebuild areas goal for an array

```
charray -rebuildareagoal 3 array1
```

The resulting output:

No feedback

An invocation example for changing the rebuild areas goal for an array

```
charray -slowwritepriority redundancy 0
```

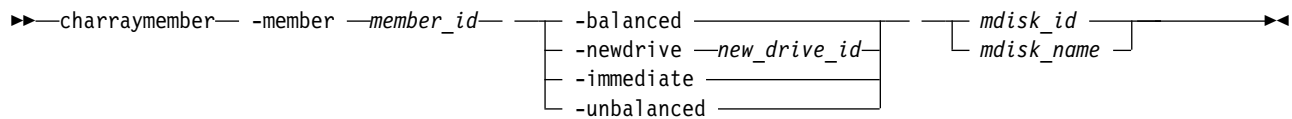
The resulting output:

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



Parameters

-member *member_id*

Identifies the array member index.

-balanced

(Optional) Forces the array member spare goals to be set to the:

- Present array member goals
- Existing exchange goals
- The newDrive goals

Note: If **-balanced** is specified and the goal of the associated array MDisk changes, the tier of the array MDisk is updated to match the new goal.

-newdrive *new_drive_id*

(Optional) Identifies the drive to add to the array.

For distributed arrays:

- If the **-newdrive** parameter is specified with the **-immediate** parameter, and the member is not associated with a rebuild area, the command begins a distributed rebuild to a rebuild area in such a way that a copyback begins immediately when the rebuild finishes.
- If the **-newdrive** parameter is specified and the member is already associated with a rebuild area, the array configures itself to use the new member, and might begin a copyback. (This occurs whether or not the **-immediate** parameter is specified.)
- If the **-immediate** parameter is not set and the **-newdrive** parameter is set (but the array member is not allocated to a rebuild area), the command fails. If the **-balanced** parameter is set, the command fails.

-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.

-unbalanced

(Optional) Forces the array member to change if the newDrive does not meet array member goals.

mdisk_id

(Either the ID or the name is required) Identifies which ID array the MDisk command applies to.

mdisk_name

(Either the ID or the name is required) Identifies which name 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.

Specify **-balanced** and the system examines the mirrored pair containing the member (including the new member drive's properties). If the array is mirror-based and the new drive is:

- On the same chain as the other member of this pair, it removes the chain-balancing goal from the array

- Not on the same chain as the other member of this pair (and there is only one mirrored pair) the array becomes chain-balanced

Because **charraymember** is member-focussed this command only operates locally to the member being operated on in terms of interacting with the new chain balanced goal.

Table 13 shows the command combination options.

Table 13. charaymember combination options

Option	Description
-balanced	<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 drive_id	<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 drive_id -balanced	The command processes the exchange and updates the member goals to the properties of the new drive.
-newdrive drive_id -unbalanced	<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 is always immediate. • Later, if drives are a sufficient member goal match, the array rebalance selects those drives. • A balancing exchange restarts the member goals.

An invocation example to swap a spare or candidate drive for a member 0 drive using exchange

```
charraymember -member 0 -newdrive 4 mdisk2
```

The resulting output:

No feedback

An invocation example to swap a spare or candidate drive for a member 1 drive and start component rebuild for the new member

```
charraymember -member 1 -newdrive 3 -immediate mdisk3
```

The resulting output:

No feedback

An invocation example to swap in a spare or candidate drive for member index 2

If there is a drive present the exchange occurs:

```
charraymember -member 2 -newdrive 4 mdisk4
```

The resulting output:

No feedback

An invocation example to force member 4 to change its spare goals to its associated drive

```
charraymember -member 4 -balanced mdisk6
```

The resulting output:

No feedback

An invocation example to force an exchange and make the array change its goals to the new drive

```
charraymember -member 3 -newdrive 9 -balanced mdisk5
```

The resulting output:

No feedback

An invocation example to force an unbalancing exchange when drive 8 does not match the goals

```
charraymember -member 2 -newdrive 8 -unbalanced mdisk5
```

The resulting output:

No feedback

An invocation example to force an immediate exchange and make the array change its goals to the new drive

```
charraymember -member 3 -newdrive 9 -balanced -immediate mdisk5
```

The resulting output:

No feedback

An invocation example to change member 24 for new drive 15 using a distributed rebuild to a rebuild area

```
svctask chararraymember -member 24 -newdrive 15 -immediate 0
```

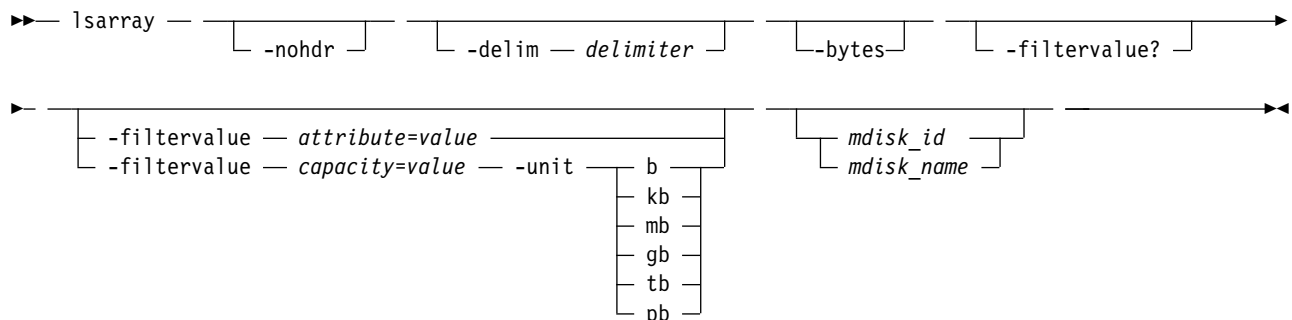
The resulting output:

No feedback

lsarray

Use the **lsarray** command to list the array MDisk.

Syntax



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) 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. In a detailed view, the data is separated from its header by the specified delimiter.

-bytes

(Optional) Requests output of capacities in bytes (instead of rounded values).

-filtervalue?

(Optional) Includes all of the valid filter attributes in the report. The following filter attributes are valid for the **lsarray** command:

- balanced
- capacity
- distributed
- mdisk_id
- mdisk_name
- mode
- mdisk_grp_id
- mdisk_grp_name
- fast_write_state
- raid_status
- raid_level
- redundancy
- spare_goal
- spare_protection_min
- status
- strip_size
- tier
- easy_tier_load

Any parameters specified with the **-filtervalue?** parameter are ignored.

-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. Use the **unit** parameter to interpret the value for size or capacity.

Note: Some filters allow the use of a wildcard when entering the command. The following rules apply to the use of wildcards when using the 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 ("");


```
lsarray -filtervalue "name=md*"
```

-filtervalue *capacity=value*

(Optional) Specifies a list of one or more filter capacity (size) values matching the specified values for the **unit** parameter. Use the **unit** parameter to interpret the value for size or capacity.

-unit *b | kb | mb | gb | tb | pb*

(Optional) The units that are used when applying the **-filtervalue** capacity, where:

- *b* = bytes
- *kb* = 1,024 bytes
- *mb* = 1,048,576 bytes
- *gb* = 1,073,741,824 bytes
- *tb* = 1,099,511,627,776 bytes
- *pb* = 1,125,899,906,842,624 bytes

mdisk_id

(Optional) The identity of the array MDisk.

mdisk_name

(Optional) The name of the array MDisk.

Description

This command returns a concise list or a detailed view of array MDisks visible to the clustered system (system). Table 14 provides the attribute values that can be displayed as output view data.

Table 14. Array output

Attribute	Values
capacity	Indicates the value for the capacity specified using the -unit parameter.
status	<ul style="list-style-type: none">• online• offline• excluded• degraded (applies only to internal MDisks)
mode	Indicates the mode. The values are: <ul style="list-style-type: none">• unmanaged• managed• image• array
quorum_index	Indicates the quorum index. The values are: <ul style="list-style-type: none">• 0• 1• 2• Blank if the MDisk is not being used as a quorum disk
block_size	Indicates the block size. The value is 512 bytes (or blank) in each block of storage.
ctrl_type	4, 6, where 6 is a flash drive attached inside a node and 4 is any other device

Table 14. Array output (continued)

Attribute	Values
raid_status	<p>Indicates the RAID status. The values are:</p> <p>offline The array is offline on all nodes.</p> <p>degraded The array has deconfigured or offline members; the array is not fully redundant.</p> <p>syncing The array members are all online, the array is synchronizing parity or mirrors to achieve redundancy.</p> <p>initializing The array members are all online, the array is initializing; the array is fully redundant.</p> <p>online The array members are all online, and the array is fully redundant.</p>
fast_write_state	<p>Indicates the cache state of the array. The values are:</p> <ul style="list-style-type: none"> • empty, which indicates that the array disk data is not changing • not_empty, which indicates that the array disk data might change • corrupt, which indicates that the array disk data is lost and the array is corrupt <p>Repair can be initiated using the recoverarray or recoverarraybysystem command.</p>
raid_level	<p>Indicates the RAID level of the array. The values are:</p> <ul style="list-style-type: none"> • RAID0 • RAID1 • RAID5 • RAID6 • RAID10
redundancy	Indicates the number of member disks that can fail concurrently without causing the array to fail.
strip_size	Indicates the strip size of the array (in KB).
spare_goal	Indicates the number of spares that the array members should be protected by. For distributed arrays, this value is blank.
spare_protection_min	Indicates the minimum number of spares that an array member is protected by. For distributed arrays, this value is blank.
balanced	<p>For nondistributed arrays, this indicates if the array is balanced to its spare goals:</p> <ul style="list-style-type: none"> • exact indicates that all populated members have the same (matching) capability and location. • yes indicates that all populated members have at least the same capability and chain, but a different enclosure or slot. • no indicates that it is unbalanced. <p>For distributed arrays, this indicates whether a superior drive class is being used for the array:</p> <ul style="list-style-type: none"> • exact indicates that the same drive class is in use. • yes indicates that at least one array member exceeds the array drive class. <p>For distributed arrays, the array must also be balanced to its rebuild area goals.</p>

Table 14. Array output (continued)

Attribute	Values
tier	<p>Indicates the tier that this array is assigned to by auto-detection (for internal arrays) or by the user:</p> <ul style="list-style-type: none"> • tier0_flash • tier1_flash • tier_enterprise • tier_nearline <p>Note: Use the chmdisk command to change this value.</p>
replacement_date	Indicates the date of a potential array failure. The format must be YYYYMMDD.
easy_tier_load	<p>Indicates the value for Easy Tier settings, and is either blank (for arrays) or one of the following values (for MDisks):</p> <ul style="list-style-type: none"> • low • medium • high • very_high
slow_write_priority	<p>Indicates the response time goal:</p> <ul style="list-style-type: none"> • latency indicates that the array is taken out of synchronization to quickly complete write operations that take excessive time. • redundancy indicates slow write operations are completed in normal time and the arrays remain synchronized.
site_id	Indicates the site value for the storage pool. This numeric value is 1, 2, 3, or blank.
site_name	Indicates the site name for the storage pool. This value is alphanumeric or is blank.
fabric_type	<p>Indicates a Fibre Channel (FC), SAS, or another type of array.</p> <ul style="list-style-type: none"> • fc indicates an array from an FC controller • sas_direct indicates an array from an SAS direct-attached controller
encrypt	<p>Indicates whether the data stored on the array is encrypted or not encrypted. The possible values are:</p> <ul style="list-style-type: none"> • yes • no
distributed	Indicates whether the array is distributed. The values are yes or no.
drive_class_id	Indicates the drive class that makes up this array. If -allowsuperior was used during array creation, the lowest used drive class ID is displayed. This value is blank for nondistributed arrays.
drive_count	Indicates the total width of the array, including rebuild areas. The value is a number from 4 to 128. The minimum value for RAID-6 and RAID-10 arrays is 6.
stripe_width	<p>Indicates the width of a single unit of redundancy within a distributed set of drives. The values are:</p> <ul style="list-style-type: none"> • Any number from 3 - 16 for RAID-5 arrays • Any number from 4 - 16 for RAID-6 arrays • An even number from 2 - 16 for RAID-10 arrays
rebuild_areas_total	Indicates the total number of rebuild areas set when the array is created. These rebuild areas provide performance but no capacity. The value is 1 - 4 for distributed array RAID-5 and distributed array RAID-6. The value is blank for nondistributed arrays.
rebuild_areas_available	Indicates the number of remaining rebuild areas within the set of drives. The value is 1 - 4 for distributed array RAID-5 and distributed array RAID-6. The value is blank for nondistributed arrays.

Table 14. Array output (continued)

Attribute	Values
rebuild_areas_goal	Indicates the rebuild areas threshold (minimum limit) at which point the array logs an error. The value is 1 - 4 for distributed array RAID-5 and distributed array RAID-6. The value is blank for nondistributed arrays.

This list defines the status fields:

online The MDisk is online and available.

degraded

(Internal MDisk 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 system.

offline

All paths to the MDisk are lost.

excluded

The MDisk is excluded from use by the system; the MDisk port error count exceeded the threshold.

A concise invocation example

```
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:encrypt
:distributed
1::online:0:mdiskgrp0:68.4GB:online:raid0:0:256:enterprise:no:yes
2:mdisk2:online:0:mdiskgrp0:88.4GB:syncing:raid5:1:256:nearline:no:no
533:mdisk533:degraded:1:mdiskgrp1:78.2GB:syncing:raid6:2:128:ssd:yes:yes
534:mdisk534:online:2:mdiskgrp1:94.2GB:inittng:raid6:2:64:ssd:yes:no
```

A detailed invocation example

```
lsarray
```

The resulting output:

```
mdisk_id 144
mdisk_name draid6_5
status online
mode array
mdisk_grp_id 1
mdisk_grp_name pool_512
capacity 5.6TB
quorum_index
block_size
controller_name
ctrl_type
ctrl_WWNN
controller_id
path_count
max_path_count
ctrl_LUN_#
```

```

UID
preferred_WWPN
active_WWPN
fast_write_state not_empty
raid_status initting
raid_level raid6
redundancy 2
strip_size 256
spare_goal
spare_protection_min
balanced exact

tier tier0 flash
replacement_date 121110090907
slow_write_priority redundancy
fabric_type
site_id
site_name
easy_tier_load
encrypt no
distributed yes
drive_class_id 1
drive_count 28
stripe_width 15
rebuild_areas_total 4
rebuild_areas_available 4
rebuild_areas_goal 2

```

A detailed invocation example

```
lsarray 1
```

The resulting output:

```

mdisk_id:1
mdisk_name:
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:
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 tier1 flash
replacement_date 121110090907
slow_write_priority:latency
site_id:3
site_name:Quorum

```

```

fabric_type:
encrypt:yes
distributed no
drive_class_id
drive_count 8
stripe_width 4
total_rebuild_areas
available_rebuild_areas
rebuild_areas_goal

```

lsarrayinitprogress

Use the **lsarrayinitprogress** command to view the progress of array background initialization that occurs after creation.

Syntax

```

➤➤ lsarrayinitprogress [-nohdr] [-filtervalue attribute_value]
➤ [-filtervalue?] [-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.

-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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):


```
lsarraysyncprogress -filtervalue mdisk_id="1*"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- estimated_completion_time
- mdisk_id
- mdisk_name
- progress

-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

(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 15 shows possible outputs.

Table 15. lsarrayinitprogress output

Attribute	Value
progress	The percentage of initialization task that has been completed.
estimated_completion_time	The expected initialization task completion time, in YYMMDDHHMMSS format.

A concise invocation example

```
lsarrayinitprogress -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

```
lsarrayinitprogress -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

```
lsarrayinitprogress -delim : mdisk4
```

The resulting output:

```
mdisk_id:mdisk_name:progress:estimated_completion_time
4:mdisk4:100:
```

lsarraylba

Use the **lsarraylba** command to permit an array logical block address (LBA) to be found from a drive and LBA.

Syntax

```
►► lsarraylba [ -nohdr ] [ -delim delimiter ]
► -drive1ba lba -drive drive_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.

-drive lba

The LBA on the drive to convert to the array LBA. The *lba* value must be specified in hex, with a 0x prefix.

-drive drive_id

The ID of the drive to view.

Description

This command permits an array LBA to be found on a drive and LBA.

The system provides volumes that have LBAs for 512-byte block sizes; however, back-end disks that have a block size of either 512 or 4096 bytes can also be used. Drives are listed in their physical size.

Use the **lsdrive** command to display the drive block size, and use the **lsdrive** or **lsarray** command to list each object (the drive and the MDisk).

Table 16 shows possible outputs.

Table 16. *lsarraylba* output

Attribute	Value
type	The type of MDisk extent allocation: <ul style="list-style-type: none">• allocated• unallocated For distributed arrays only: <ul style="list-style-type: none">• If the LBA is an unused rebuild area, this value displays <code>rebuild_area</code>.• If the LBA is a used rebuild area, this value displays <code>allocated</code>.
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).

An invocation example to map drive 2 LBA -xff to MDisk 2 LBA 0xff

```
lsarraylba -delim : -drive1ba 0xff -drive 2
```

The resulting output:

```
mdisk_id:mdisk_name:type:mdisk_lba:mdisk_start:mdisk_end:drive_start:drive_end
0:mdisk2:allocated:0x00000000000001ff:0x0000000000000100:0x00000000000001ff:
0x0000000000000000:0x00000000000000ff
```

An invocation example for an allocated space

```
lsarraylba -drive1ba 0x00 -drive 2
```

The resulting output:

mdisk_id	mdisk_name	type	mdisk_lba	mdisk_start	mdisk_end	drive_start	drive_end
1	mdisk1	allocated	0x0000000000000000	0x0000000000000000	0x00000000000001FF	0x0000000000000000	0x0000000000000000

An invocation example for an unused rebuild area

```
lsarraylba -drive1ba 0x00 -drive 16
```

The resulting output:

mdisk_id	mdisk_name	type	mdisk_lba	mdisk_start	mdisk_end	drive_start	drive_end
3	mdisk3	rebuild_area				0x0000000000000000	0x0000000000000000

Isarraymember

Use the **lsarraymember** command to list the member drives of one or more array MDisks.

Syntax

```
lsarraymember - [-nohdr] [-filtervalue attribute=value]
                [-filtervalue?] [-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.

-filtervalue attribute=value

(Optional) Specifies a list of one or more filter attributes matching the specified values; see **-filtervalue?** for the supported attributes.

Note: Some filters allow the use of a wildcard when entering the command. The following rules apply to the use of wildcards when using the 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 (" "):

```
lsarraymember -filtervalue "mdisk_name=md*"
```

-filtervalue?

(Optional) Includes all of the valid filter attributes in the report. The following filter attributes are valid for the **lsarraymember** command:

- `mdisk_id`
- `mdisk_name`
- `member_id`
- `drive_id`
- `new_drive_id`
- `spare_protection`
- `balanced`

Any parameters specified with the **-filtervalue?** parameter are ignored.

-delim delimiter

(Optional) By default, in a concise view all columns of data are space-separated, with the width of each column set to the maximum possible width of each item of data. In a detailed view, each item of data is an individual row, and if displaying headers, 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. Enter **-delim :** on the command line, and 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 specified *delimiter* separates the data from its header.

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 unoccupied by a drive. The positions determine how mirroring the RAID's takes place. For example, determining if *x* is mirrored to *y* for RAID-10, where parity starts from RAID-5, which is for enclosure-based systems only.

Table 17 shows the potential output for this command.

Table 17. lsarraymember output

Attribute	Value
<code>member_id</code>	Specifies the identity of the array member. It represents drive order in RAID array
<code>drive_id</code>	Specifies the identity of the drive for member ID, or the source drive if an exchange is in progress. It is blank if there is no drive configured.
<code>new_drive_id</code>	Specifies the ID of the drive exchanged with this member ID. It is blank if there is no ID.
<code>spare_protection</code>	Specifies the number of non-degrading spares for the member. This includes spare drives with different attributes from the array member goals that perform equally or better than the array member goals. For distributed array members this field is blank.

Table 17. *lsarraymember* output (continued)

Attribute	Value
balanced	<p>For nondistributed arrays, this describes if the array is balanced to its spare goals:</p> <ul style="list-style-type: none"> • exact indicates that all populated members have the same (matching) capability and location. • yes indicates that all populated members have at least the same capability and chain, but a different enclosure or slot. • no indicates that it is unbalanced. <p>For distributed arrays, this indicates whether a superior drive class is being used for the array:</p> <ul style="list-style-type: none"> • exact indicates that the same drive class is in use. • yes indicates that the drive exceeds the array drive class. <p>For distributed arrays, the array must also be balanced to its rebuild area goals.</p>
slow_write_count	Indicates the number of times this member has become unsynchronized because of high response time on write I/O operations.
slow_write_time_last	Creates a timestamp of when the component last became unsynchronized. The time format is YYMMDDhhmmss in clustered system time. No time is indicated if the value is slow_write_count is 0.

A concise invocation example

```
lsarraymember -delim :
```

The resulting output:

```
lsarraymember -delim :
mdisk_id:mdisk_name:member_id:drive_id:new_drive_id:spare_protection
      :balanced:slow_write_count:slow_write_time_last
2:mdisk1:0:55::1:exact:4:130103202158
2:mdisk1:1:56::1:exact:1:130103203930
2:mdisk2:0:0::2:exact:0:
2:mdisk2:1:2:5:3:exact:2:130103204044
2:mdisk2:2::::::
2:mdisk2:3:8::0:no::
```

A concise invocation example (qualified with MDisk)

```
lsarraymember -delim : mdisk_2
```

The resulting output:

```
tmdisk_id:mdisk_name:member_id:drive_id:new_drive_id:spare_protection
      :balanced:slow_write_count:slow_write_time_last
2:mdisk2:0:0::2:exact:4:130103202158
2:mdisk2:1:2:5:3:exact:1:130103203930
2:mdisk2:2::::::0:
2:mdisk2:3:8::0:no:2:130103204044
```

Note: From this output, you can see that:

- The array has four members (possibly a 4-member RAID-10 array). You cannot use RAID-10 with distributed arrays.
- The second array member is undergoing exchange for drive5.
- The third array member is not configured. It might be 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)

```
lsarraymember -delim :
```

The resulting output:

```
mdisk_id:mdisk_name:member_id:drive_id:new_drive_id:spare_protection
      :balanced:slow_write_count:slow_write_time_last
2:mdisk1:0:55:::1:exact:4:130103202158
2:mdisk1:1:56:::1:exact:1:130103203930
2:mdisk2:0:0:::2:exact:0:
2:mdisk2:1:2:5::3:exact:2:130103204044
2:mdisk2:2:::0:
2:mdisk2:3:8:::0:no::
```

An invocation example (an array with a change in membership from (55,56) to (55,57,58))

```
lsarraymember -delim : mdisk_3
```

The resulting output:

```
mdisk_id:mdisk_name:member_id:drive_id:new_drive_id:spare_protection
      :balanced:slow_write_count:slow_write_time_last
3:mdisk3:0:55::55:1:exact:4:130103202158
3:mdisk3:1:56::57:1:exact:1:130103203930
3:mdisk3:2:::58:1:exact:0:
```

An invocation example (an array with a change in membership from (55,57,58) to (55,56))

```
lsarraymember -delim : mdisk_3
```

The resulting output:

```
mdisk_id:mdisk_name:member_id:drive_id:new_drive_id:spare_protection
      :balanced:slow_write_count:slow_write_time_last
3:mdisk3:0:55::55:1:exact:4:130103202158
3:mdisk3:1:57::56:1:exact:1:130103203930
3:mdisk3:2:58:::1:exact:0:
```

lsarraymembergoals

Use the **lsarraymembergoals** command to list the spare goals for member drives of one or more array MDisks.

Syntax

```
➤➤ lsarraymembergoals [ -filtervalue — attribute_value ] [ -filtervalue? ]
➤ [ -delim — delimiter ] [ -bytes ] [ mdisk_id ] [ mdisk_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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):


```
lsarraymembergoals -filtervalue mdisk_id="1*"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- capacity_goal
- drive_id
- enclosure_id_goal
- estimated_completion_time
- mdisk_id
- mdisk_name
- member_id
- node_id_goal
- progress
- RPM_goal
- slot_id_goal
- tech_type_goal
- drive_class_id_goal

-delim delimiter

(Optional) By default, in a concise view all columns of data are space-separated, with the width of each column set to the maximum possible width of each item of data. In a detailed view, each item of data is an individual row, and if displaying headers, 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. Enter **-delim :** on the command line, and 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* .

-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 lists the spare goals for member drives of one or more array MDisks. Table 18 provides the potential output for this command.

Table 18. lsarraymembergoals output

Attribute	Values
member_id	Indicates the ID of the array member which represents the drive order in the RAID array.
drive_id	Indicates the ID of the drive for the member ID (it is blank if none is configured).
capacity_goal	Indicates the capacity goal for the array member (it is the same for all members in the array).

Table 18. *lsarraymembergoals* output (continued)

Attribute	Values
tech_type_goal	Indicates the technology goal for the array member: <ul style="list-style-type: none"> • tier0_flash • tier1_flash • tier_enterprise • tier_nearline
RPM_goal	Indicates the drive RPM characteristic that the array member should have (it is blank for flash drives).
enclosure_id_goal	Indicates the ID of the enclosure that should contain the array member (it is blank if any enclosure can be selected).
slot_id_goal	Indicates the ID of the slot in the enclosure that should contain the array member.
node_id_goal	Indicates the ID of the node that should contain this array member.
enclosure_balance_goal	Indicates if a member drive's enclosure must be connected to the same SAS port as the drive that set the array goals.
block_size_goal	Indicates the array member block size. The value is either 512 or 4096. This value is the same for all member drives in the array, and is the smallest value for the block size of one of the original drives or the set of drives in the array when it is set to be balanced.
drive_class_id_goal	Indicates the preferred drive class for this array member (the value is blank for nondistributed arrays).

An invocation example (a four-member RAID-10 SAS array that is split across chains)

You cannot use RAID-10 with distributed arrays.

```
lsarraymembergoals -delim : mdisk_2
```

The resulting output:

```
mdisk_id:mdisk_name:member_id:drive_id:capacity_goal:
tech_type_goal:RPM_goal:enclosure_id_goal:slot_id_goal:enclosure_balance_goal:node_id_goal:block_size_goal:drive_class_id_goal:
2:mdisk2:0:0:68.4GB:tier0_flash:15000:1:1:no:512:0
2:mdisk2:1:17:68.4GB:tier0_flash:15000:1:2:no:512:0
2:mdisk2:2:1:68.4GB:tier0_flash:15000:14:1:no:512:2
2:mdisk2:3:18:68.4GB:tier0_flash:15000:14:2:no:512:2
```

An invocation example

```
lsarraymembergoals -filtervalue block_size_goal=4096
```

The resulting output:

mdisk_id	mdisk_name	member_id	drive_id	capacity_goal	tech_type_goal	RPM_goal	enclosure_id_goal	slot_id_goal	node_id_goal	enclosure_balance_goal	block_size_goal	drive_class_id_goal
4	r10_array	0	43	1.6TB	tier0_flash	10000	1	21		no	4096	0
4	r10_array	1	44	1.6TB	tier0_flash	10000	1	18		no	4096	0
4	r10_array	2	45	1.6TB	tier0_flash	10000	1	20		no	40962	
4	r10_array	3	46	1.6TB	tier0_flash	10000	2	5		no	40962	

```
lsarraymembergoals
```

The resulting output:


mdisk_id	mdisk_name	member_id	drive_id	capacity_goal	tech_type_goal	RPM_goal	enclosure_id_goal	slot_id_goal	node_id_goal	enclosure_balance_goal	block_size_goal	drive_class_id_goal
0	r10_array	0	1	278.9GB	tier1_flash	15000	1	2		no		
0	r10_array	1	10	278.9GB	tier1_flash	15000	1	3		no		
0	r10_array	2	9	278.9GB	tier1_flash	15000	1	4		no		
0	r10_array	3	0	278.9GB	tier1_flash	15000	1	5		no		
0	r10_array	4	6	278.9GB	tier1_flash	15000	1	6		no		
0	r10_array	5	7	278.9GB	tier1_flash	15000	1	7		no		
0	r10_array	6	18	278.9GB	tier1_flash	15000	1	8		no		

0	r10_array	7	21	278.9GB	tier1_flash	15000	1	9	no
1	r0_array	0	15	278.9GB	tier1_flash	15000	1	10	no
1	r0_array	1	22	278.9GB	tier1_flash	15000	1	11	no
1	r0_array	2	13	278.9GB	tier1_flash	15000	1	12	no
1	r0_array	3	5	278.9GB	tier1_flash	15000	1	13	no
2	r1_array3	0	8	278.9GB	tier1_flash	15000	1	14	no
2	r1_array3	1	4	278.9GB	tier1_flash	15000	1	15	no
3	r1_array1	0	16	278.9GB	tier1_flash	15000	1	16	no
3	r1_array1	1	12	278.9GB	tier1_flash	15000	1	17	no
4	r1_array2	0	17	278.9GB	tier1_flash	15000	1	20	no
4	r1_array2	1	19	278.9GB	tier1_flash	15000	1	19	no

Isarraymemberprogress

Use the **lsarraymemberprogress** command to display array member background process status.

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.

-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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):


```
lsarraymemberprogress -filtervalue mdisk_id="1*"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- estimated_completion_time
- drive_id
- mdisk_id
- mdisk_name
- member_id
- new_drive_id
- progress
- task

-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

(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 19 provides the potential output for this command.

Table 19. Isarraymemberprogress output

Attribute	Value
member_id	Indicates the array member index.
drive_id	Indicates the ID of the drive.
task	Indicates the identity of task being performed by the array member: <ul style="list-style-type: none"> • rebuild indicates the array is recovering all the data on the component (after it was removed) • exchange indicates the component is copying data to another drive • resync indicates this member is unsynchronized and is performing write operations that were completed early • copyback indicates this member is copying data to an array member that has recently become active. This applies to distributed arrays. Note: For example, if the drive fails the array is rebuilt. If the drive does not fail or is replaced, a copyback occurs to write data back to the array member.
new_drive_id	The identity of drive being exchanged.
progress	Indicates the task percentage complete.
estimated_completion_time	Indicates the expected task completion time in the format YYMMDDHHMMSS. It is blank if completion time is unknown.

A concise invocation example

```
lsarraymemberprogress -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:
```



```
lsarraymemberprogress mdisk_2
```

```
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:
```

1sarraymemberprogress

mdisk_id	mdisk_name	member_id	drive_id	task	new_drive_id	progress	estimated_completion_time
3	mdisk3	5	1	resync		95	121203193637
3	mdisk3	6	2	rebuild		0	121203234321
3	mdisk3	7	3	exchange	18	0	121204033229

1sarraymemberprogress

mdisk_id	mdisk_name	member_id	drive_id	task	new_drive_id	progress	estimated_completion_time
3	mdisk3	4	7	copyback		10	150710165446

Use the **lsarrayrecommendation** command to view a recommended configuration for the specified drive class and number of drives.

```

▶▶ — lsarrayrecommendation — — -driveclass —drive_class_id_list— —————▶
▶ -drivecount —drive_count_list— ┌ —nohdr┐ ┌ —filtervalue?┐ —————▶
                                   └—————┘ └—————┘
▶ ┌ —filtervalue — attribute=value ┐ ┌ —delim — delimiter ┐ ┌ mdiskgrp_id ┐ —▶
  └—————┘ └—————┘ └—————┘ └mdiskgrp name┘

```

-driveclass *drive class id list*

(Required) Specifies the drive class, or classes, for which the array recommendation is made. You must specify at least 1 *drive_class_id_list* value. You can specify a total of 32 *drive_class_id_list* values on a single command; however, you must separate each value with a colon character (:).

-drivecount *drive count list*

(Required) Specifies the number of drives for which to make recommendation. You must specify at least 1 *drive_count* value. You can specify a total of 32 *drive_count* values on a single command; however, you must separate each value with a colon character (:).

Remember: Each drive_count value must be a numerical value between 2 and 128 (only redundant arrays are considered).

-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.

-filtervalue attribute=value

(Optional) Specifies a list of one or more filter attributes that match the specified values; see **-filtervalue?** for the supported attributes.

Note: Some filters allow the use of a wildcard; the following rules apply when using a wildcard character:

- 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.
- You must enclose the filter entry within double quotation marks ("").

-filtervalue?

(Optional) Includes all of the valid filter attributes in the report. The following filter attributes are valid for the **lsarrayrecommendation** command:

- raid_level
- distributed

Any parameters that are specified with the **-filtervalue?** parameter are ignored.

-delim delimiter

(Optional) By default, in a concise view all columns of data are space-separated, with the width of each column set to the maximum width of each item of data. In a detailed view, each item of data is an individual row, and if you display headers, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a 1-byte character. 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 specified *delimiter* separates the data from its header.

mdiskgrp_id

mdiskgrp_name

(Required) The ID or name of the pool for which to make the recommendation.

Description

This command displays the system-recommended array configuration for a specific drive class and number of drives.

Encrypted pools can be recommended or used if the specified storage pool is encrypted.

Table 20 provides the attribute values that can be displayed as output view data.

Table 20. lsarrayrecommendation output

Attribute	Possible Values
mdiskgrp_id	Indicates the MDisk group ID
mdiskgrp_name	Indicates the MDisk group name
drive_class_id	Indicates the drive class ID for the recommendation.
raid_level	Indicates the recommended RAID level.

Table 20. *lsarrayrecommendation* output (continued)

Attribute	Possible Values
distributed	Indicates whether this is a distributed array recommendation.
min_stripe_width	Indicates the stripe width minimum.
max_stripe_width	Indicates the stripe width maximum.
stripe_width	Indicates the recommended stripe size for this storage pool.
rebuild_areas	Indicates the recommended number of rebuild areas.
strip_size	Indicates the recommended strip size for this storage pool.
drive_count	Indicates the number of drives to configure per array.
array_count	Indicates the number of arrays that can be created at the recommended drive count.
capacity	Indicates the expected usable capacity for virtualization for the array or arrays.

Note: The default recommendation for each drive class is the first row that appears in the output for that drive class. All other rows for each drive class are not sorted by order of recommendation.

An invocation example

Making a recommendation for more rebuild areas and larger stripe width because there are other existing arrays in the MDisk group:

```
lsarrayrecommendation -driveclass 2 -drivecount 60 dist_pool
```

The detailed resulting output:

```
mdisk_grp_id mdisk_grp_name drive_class_id raid_level distributed min_stripe_width max_stripe_width stripe_width rebuild_areas strip_size
0            dist_pool      2              raid6        yes           5             16             12             3             256
```

An invocation example

Making a recommendation for a new (empty) storage pool and with a row for each raid level:

```
lsarrayrecommendation -driveclass 2 -drivecount 80 mdiskgrp0
```

The detailed resulting output:

```
mdisk_grp_id mdisk_grp_name drive_class_id raid_level distributed min_stripe_width max_stripe_width stripe_width rebuild_areas strip_size drive_count array_count capacity
0            mdiskgrp0      2              raid5        yes           3             16              6              2             256         40         2         410.1TB
0            mdiskgrp0      2              raid6        yes           5             16             12              3             256         40         2         400.1TB
0            mdiskgrp0      2              raid5        yes           2             16              8              2             256         40         2         380.1TB
0            mdiskgrp0      2              raid1        no            2             16              2             10             256         10         8         410.1TB
0            mdiskgrp0      2              raid5        no            3             16             10             7             256         10         7         410.1TB
0            mdiskgrp0      2              raid6        no            5             16             10             7             256         10         7         400.1TB
0            mdiskgrp0      2              raid5        no            2             16              8              9             256          8         9         380.1TB
```

An invocation example

Making a recommendation for multiple drive classes for a new (empty) storage pool:

```
lsarrayrecommendation -driveclass 3:5 -drivecount 80:24 mdiskgrp0
```

The detailed resulting output:


```
mdisk_grp_id mdisk_grp_name drive_class_id raid_level distributed min_stripe_width max_stripe_width stripe_width rebuild_areas strip_size drive_count array_count capacity
0            mdiskgrp0      3              raid5        yes           3             16              6              2             256         40         2         500.4TB
0            mdiskgrp0      3              raid6        yes           5             16             12              3             256         40         2         480.4TB
0            mdiskgrp0      3              raid5        yes           2             16              8              2             256         40         2         450.4TB
0            mdiskgrp0      3              raid1        no            2             16              2             10             256         10         8         400.6TB
0            mdiskgrp0      3              raid5        no            3             16             10             7             256         10         7         500.6TB
0            mdiskgrp0      3              raid6        no            5             16             10             7             256         10         7         480.6TB
```

0	mdiskgrp0	3	raid5	no	2	16	8		256	8	9	450.6TB
0	mdiskgrp0	5	raid5	yes	3	16	6	2	256	12	2	200.7TB
0	mdiskgrp0	5	raid6	yes	5	16	12	3	256	12	2	180.7TB
0	mdiskgrp0	5	raid5	yes	2	16	8	2	256	8	3	150.7TB
0	mdiskgrp0	5	raid1	no	2	16	2		256	8	3	100.2TB
0	mdiskgrp0	5	raid5	no	3	16	12		256	12	2	200.2TB
0	mdiskgrp0	5	raid6	no	5	16	12		256	12	2	180.2TB
0	mdiskgrp0	5	raid5	no	2	16	8		256	8	3	150.2TB

Isarraysyncprogress

Use the **Isarraysyncprogress** command to display how synchronized a RAID array is.

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.

-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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):


```
Isarraysyncprogress -filtervalue mdisk_id="1*"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- estimated_completion_time
- mdisk_id
- mdisk_name
- progress

-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

(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 21 provides the potential output.

Table 21. 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

```
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

```
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

```
lsarraysyncprogress -delim : mdisk_10
```

The resulting output:


```
mdisk_id:mdisk_name:progress:estimated_completion_time
0:mdisk0:100:
```

lspotentialarraysize

Use the **lspotentialarraysize** command to display the size of a potential array for a specified drive count, drive class, and RAID level in the specified MDisk group.

Syntax

```
►► lspotentialarraysize — -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, with the width of each column set to the maximum width of each item of data. In a detailed view, each item of data is an individual row, and if you display headers, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a 1-byte character. 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 specified *delimiter* separates the data from its header.

-drivecount 3 - 128

(Required) Specifies the number of the drives. The value must be a number from 3 - 128.

-driveclass drive_class_id

(Required) Specifies the drive class. The *driveclass_id* value must be a number.

-level raid1 | raid5 | raid6 | raid10

(Required) Specifies one of the following RAID levels for the array that is being created. The values are:

- *raid1*
- *raid5*
- *raid6*
- *raid10*

-stripewidth 2 - 16

(Required) Indicates the width of a single unit of redundancy within a distributed set of drives. The value must be:

- RAID-1: 2 - 16
- RAID-5: 3 - 16
- RAID-6: 5 - 16
- RAID-10: 2, 4, 6, 8, 10, 12, 14, 16 (You cannot use RAID-10 with distributed arrays.)

-rebuildareas 1 | 2 | 3 | 4

(Optional) Specifies the number of rebuild areas in the array. This value must be 1 - 4 (inclusive) for RAID-5 and RAID-6 arrays.

Note: This parameter is only applicable for distributed arrays.

-strip 128 | 256

(Optional) Specifies sets the strip size in KiB for the array that is being configured. The values are 128 or 256.

mdiskgrp_id | mdiskgrp_name

(Required) Indicates the MDisk array ID or name.

Description

This command displays the size of a potential array for a specified drive count, class, and RAID level.

You can use this command to define potential sizes for nondistributed and distributed arrays. (Distributed array descriptions are triggered by using **-rebuildareas**.) This command assists with the configuration options provided during array creation, and estimates the array capacity if it were to be configured

Table 22 provides the attribute values that can be displayed as output view data.

Table 22. lspotentialarraysize output

Attribute	Possible Values
capacity	Indicates the expected usable capacity for virtualization for the array or arrays.

An invocation example that uses a small drive count for a distributed array

```
lspotentialarraysize -driveclass 4 -drivecount 40 -level raid5 -stripewidth 6 -rebuildareas 2 mdiskgrp1
```

The detailed resulting output:

```
capacity
115.2TB
```

An invocation example that uses a different class and fewer rebuild areas

```
lspotentialarraysize -driveclass 4 -drivecount 100 -level raid5 -stripewidth 8 -strip 128 -rebuildareas 1 mdiskgrp1
```

The detailed resulting output:

```
capacity
172.4TB
```

An invocation example that uses the same class and no rebuild areas

```
lspotentialarraysize -driveclass 4 -drivecount 100 -level raid5 -stripewidth 8 -strip 128 1
```

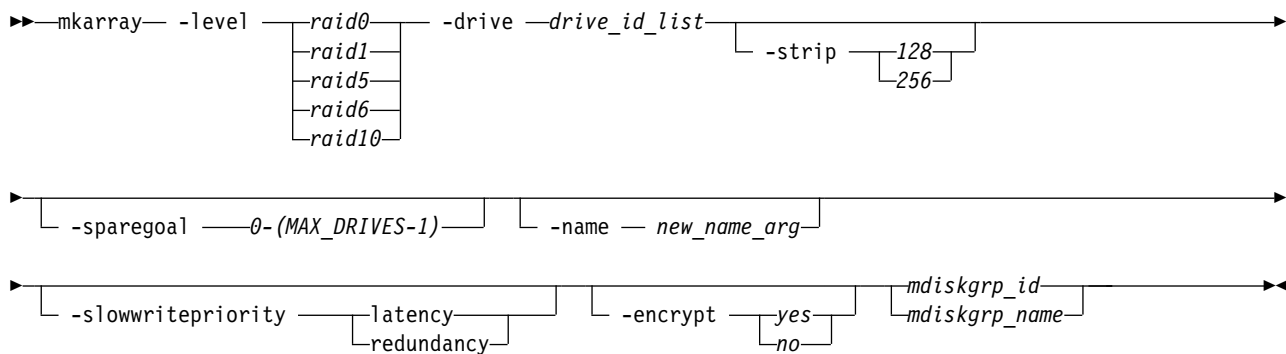
The detailed resulting output:

```
capacity
184.3TB
```

mkarray

Use the **mkarray** command to create an MDisk array and add it to a storage pool. This command applies to nondistributed arrays. (Use the **mkdistributedarray** command to create distributed arrays.)

Syntax



Parameters

-level

(Required) Sets the RAID level for the array MDisk being created.

The following requirements apply for RAID levels:

- RAID-0: Stripes data across all members, provides no redundancy.
- RAID-1: Mirrored pair of drives, allows reading from either drive. Can tolerate either drive failing.
- RAID-5: These arrays stripe data over the member drives with one parity strip on every stripe and can tolerate no more than one member drive failure.
- RAID-6: These arrays stripe data over the member drives with two parity strips on every stripe and can tolerate any two concurrent member drive failures.
- RAID-10: These arrays are in a set of up to eight mirrored pairs with the data striped across mirrors. They can tolerate the failure of one drive in each mirror and they allow reading from both drives in a mirror. (You cannot use RAID-10 with distributed arrays.)

Restriction: RAID-5 and RAID-6 are for enclosure-based systems only.

-drive drive_id_list

(Optional) Identifies the drive or drives to use as members of the RAID array.

For RAID-1 and RAID-10 arrays, drives are specified as a sequence of mirrored drive pairs. For example, if an array is created with -drive a:b:c:d, drive *b* contains the mirror copy of drive *a*, and drive *d* contains the mirror copy of drive *c*. (You cannot use RAID-10 with distributed arrays.)

This list shows how many member drives are allowed in each supported RAID type:

- RAID-0: Allows one-member to eight-member drives.

Note: Internal drives must be in the same node.

- RAID-1: Allows two-member drives.
- RAID-5, which is for enclosure-based systems only: Allows three-member to sixteen-member drives.
- RAID-6, which is for enclosure-based systems only: Allows five-member to sixteen-member drives.
- RAID-10: Allows drives with:
 - Two members
 - Four members
 - Six members
 - Eight members
 - Ten members
 - Twelve members
 - Fourteen members

- Sixteen members

Each pair of drives must contain a drive from a node in the I/O group and a drive from the other node. (You cannot use RAID-10 with distributed arrays.)

-strip *128 | 256*

(Optional) Sets strip size (in KB) for the array MDisk being created. The default is 256 KB.

-sparegoal *0-(MAX_DRIVES-1)*

(Optional) Sets the number of spares that this array's members must be protected by. The default is 1 (except for RAID-0 arrays, which have a default of 0).

-name *new_name_arg*

(Optional) Specifies the name to which you want to apply the array MDisk.

-slowwritepriority *latency | redundancy*

(Optional) Controls array ability to complete write operations that take too long, even if it temporarily compromises redundancy.

The value can be either *latency* or *redundancy*:

- *latency* implies the feature is enabled for normal I/O operations
- *redundancy* implies the feature is not enabled for normal I/O operations

The default value is *latency* mode for existing arrays, unless the array is RAID-0 (in which case *redundancy* mode is required).

Important: Do not change the mode of a RAID-0 array.

-encrypt *yes | no*

(Optional) Specifies the array to encrypt. The values are *yes* and *no*.

This parameter defaults to *yes* when **lscryption** has its status set to enabled and all nodes in the I/O group that the array is being defined on are encryption-capable.

Note: The value can be *yes* only if encryption is enabled on the array's I/O group.

mdiskgrp_id | mdiskgrp_name

(Required) Identifies the storage pool (by name or ID) to which you want to add the created array MDisk.

Description

This command creates an array MDisk RAID array and adds it to a storage pool. Although the array tier is automatically determined, you can change it later using the **chmdisk** command.

An array MDisk being added to a storage pool used for active-active relationships must match other MDisk in the storage pool.

Remember: This command cannot be used to add an array to a child pool.

If the `raid_level` is RAID-1 or RAID-10, and the drive list contains drives that do not share a SAS port connection chain, the array attempts to continue to maintain the location balance between the mirrored pairs. (You cannot use RAID-10 with distributed arrays.) Configuration changes indicate a member drive might not be goal-balanced depending on its current chain. This is relative to both the drive that created the array member goals and the current chain of the mirror partner.

If the MDisk group has an encryption key, the array must be encrypted.

An invocation example (to create arrays)

```
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)

```
mkarray -level raid1 -drive 4:5 -strip 128 mdiskgrp_4
```

The resulting output:

```
MDisk, id [1], successfully created
```

An invocation example for creating an unencrypted array on encrypted hardware

```
mkarray -level raid10 -drives 0:1:2:3:4:5 -encrypt no 0
```

The resulting output:

```
MDisk, id [1], successfully created
```

mkdistributedarray

Use the **mkdistributedarray** command to create a distributed array and add it to a storage pool. (Use the **mkarray** command to create nondistributed arrays.)

Syntax

```
►►— mkdistributedarray — -level ——— raid5 ——— -driveclass ——— driveclass_id ———►
                               └── raid6 ───┘

►— -drivecount — 4 - 128 ———┬── -stripwidth — 3-16 ───┬── -allowsuperior ───►
                               └──────────────────┘      └──────────────────┘

►┬── -rebuildareas ┬── 1 ───┬── -rebuildareasgoal ┬── 0 ───┬── -strip ┬── 128 ───►
  └── 2 ───┘         └── 1 ───┘                   └── 2 ───┘         └── 256 ───┘
  └── 3 ───┘         └── 2 ───┘                   └── 3 ───┘
  └── 4 ───┘         └── 3 ───┘                   └── 4 ───┘
                     └── 4 ───┘

►┬── -name — new_name_arg ───┬── -encrypt ┬── yes ───►
  └──────────────────┘             └── no ───┘

►┬── -slowwritepriority ┬── latency ───┬── mdiskgrp_id ───►
  └── redundancy ───┘             └── mdiskgrp_name ───┘
```

Parameters

-level *raid5* | *raid6*

(Required) Specifies the RAID level for the array being created. The values are:

- *raid5*
- *raid6*

-driveclass *driveclass_id*

(Required) Specifies the class being used to create the array. The *driveclass_id* must be a numeric value (specified with the **lsdriveclass** command).

-drivecount 4 - 128

(Required) Specifies the number of drives to use for the array. The minimum drive count for:

- RAID-5: 4
- RAID-6: 6

-stripewidth 3-16

(Optional) Indicates the width of a single unit of redundancy within a distributed set of drives. The value must be:

- RAID-5: 3 - 16
- RAID-6: 5 - 16

-allowsuperior

(Optional) Specifies that you can use drives that are not an exact match to the drive class used when creating the array (such as drives that use different capacity or technology). The system attempts to select the closest match to the class when satisfying the drive count. You can select higher capacity members of the same technology type before you select higher technology members.

Note: For a drive A to be considered superior to drive B, these situations must be true:

1. Drives A and B are use=candidate
2. Drives A and B are in the same I/O group.
3. Drive A's speed (RPM) is equal to or greater than drive B's. Solid-state drives (SSDs) are higher speed than all hard disk drives (HDDs).
4. Drive A's capacity is equal to or greater than drive B's.
5. Drive A has a block size that is smaller than or equal to drive B.

-rebuildareas 1 | 2 | 3 | 4

(Optional) Specifies the reserved capacity that is distributed across all drives available to an array. This capacity restores data after a drive failure. The values are:

- 1
- 2
- 3
- 4

The value is 1 - 4 (inclusive) for RAID-5 and RAID-6 arrays.

-rebuildareasgoal 0 | 1 | 2 | 3 | 4

(Optional) Specifies the number of rebuild areas that the array can target to keep available. If the number available in the array falls below this number, a system alert is raised.

Note: The goal value should not exceed the number of rebuild areas specified for the array. The values are:

- 0
- 1
- 2
- 3
- 4

-strip 128 | 256

(Optional) Specifies the strip size in KiB for the array being configured. The values are 128 or 256.

-name *new_name_arg*

(Optional) Specifies the name of the array.

-encrypt *yes* | *no*

(Optional) Specifies the array to encrypt. The values are *yes* and *no*. This parameter defaults to *yes* when **lencryption** has its status set to enabled and all nodes in the I/O group that the array is being defined on are encryption-capable.

Note: The value can be *yes* only if encryption is enabled on the array's I/O group.

If you specify **-encrypt yes** when the I/O group does not support encryption, the command fails.

-slowwritepriority *latency* | *redundancy*

(Optional) Controls array ability to complete write operations that take too long, even if it temporarily compromises redundancy.

The value can be either *latency* or *redundancy*:

- *latency* implies the feature is enabled for normal I/O operations
- *redundancy* implies the feature is not enabled for normal I/O operations

The default value is *latency* mode for existing arrays).

Important: An array can cause member drives to become unsynchronized (to preserve response time) if the value is *latency*. If the value is *redundancy*, the array cannot cause member drives to become unsynchronized (to preserve time) and I/O performance is impacted.

mdiskgrp_id | *mdiskgrp_name*

(Required) Indicates the MDisk array ID or name.

Description

This command creates distributed arrays.

Remember: You cannot create an unencrypted array to add to an encrypted storage pool.

Each distributed array occupies 16 slots, which starts at an MDisk ID that is divisible by 16. See the **lsmdisk** command for more information.

An invocation example to create an array that uses 40 drives of class 3 with 3 rebuild areas

```
mkdistributedarray -level raid6 -driveclass 3 -drivecount 40 -stripewidth 10 -rebuildareas 3 mdiskgrp5
```

The detailed resulting output:

```
MDisk, id [16], successfully created
```

An invocation example to create an array with a drive class

```
mkdistributedarray -level raid5 -driveclass 0 -drivecount 56 -stripewidth 8 -allowsuperior mdiskgrp2
```

The detailed resulting output:

```
MDisk, id [32], successfully created
```

An invocation example to create an array with maximum rebuild areas that logs an error on using the second rebuild area

```
mkdistributedarray -level raid5 -driveclass 5 -drivecount 60 -rebuildareas 4 -rebuildareasgoal 3 mdiskgrp2
```

The detailed resulting output:

```
MDisk, id [16], successfully created
```

An invocation example to create an array that might affect I/O performance

```
mkdistributedarray -driveclass 10 -slowwritepriority redundancy 0
```

The detailed resulting output:

```
MDisk, id [32], successfully created
```

An invocation example to make an encrypted distributed array that uses 40 drives of class 3 with 3 rebuild areas

```
mkdistributedarray -level raid6 -driveclass 3 -drivecount 40 -stripewidth 10 -rebuildareas 3 -encrypt yes mdiskgrp5
```

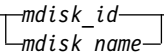
The detailed resulting output:

```
MDisk, id [16], successfully created
```

recoverarray

Use the **recoverarray** command to recover a specific corrupt array in a dead domain scenario.

Syntax

```
►►—recoverarray——————►►
```

Parameters

mdisk_id

(Optional) Identifies (by ID) the specific array to recover.

mdisk_name

(Optional) Identifies (by user-assigned name) the specific array to recover.

Description

This command recovers a specific corrupt array. An array has metadata representing ongoing or pending platform writes, which are lost when the domain nodes are lost.

An invocation example

```
recoverarray mdisk_1
```

The resulting output:

There is no output if the command is successful.

recoverarraybycluster (Discontinued)

Attention: The **recoverarraybycluster** command has been discontinued. Use the **recoverarraybysystem** command instead.

recoverarraybysystem

Use the **recoverarraybysystem** command to recover corrupt arrays in a dead domain scenario.

Syntax

```
►►—recoverarraybysystem—————►►
```

Parameters

None.

Description

Use the **recoverarraybysystem** command to recover corrupt arrays in a dead domain scenario.

An invocation example

```
recoverarraybysystem
```

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

```
►►—rmarray— -mdisk — mdisk_id_list | mdisk_name_list — mdiskgrp_id | mdiskgrp_name —◄◄  
                        └──────────┘ └── -force ─┘ └──────────┘
```

Parameters

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Required) Identifies the array MDisk or a colon-delimited list of MDisks to remove from the storage pool.

-force

(Optional) Forces a remove when the MDisk has allocated extents by migrating the used extents to free extents in the storage pool.

mdiskgrp_id | *mdiskgrp_name*

(Required) Identifies (by name or ID) the storage pool to remove the created array MDisk from.

Description

This command removes an array MDisk from the configuration. Each array is divided into candidate drives.

Remember: This command cannot be used to remove an array MDisk from a child pool.

An invocation example

```
rmarray -mdisk 6 mdiskgrp_10
```

The resulting output:

No feedback

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:

- **dumpconfig**
- **cpdumps**
- **finderr**
- **dumperrlog**

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

```
➤ catauditlog [ -nohdr ] [ -delim delimiter ] [ -first number_of_entries_to_return ] ➤
```

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.

-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.

Use this command to display the in-memory audit log. Use the **dumppauditlog** command to manually dump the contents of the in-memory audit log to a file on the current configuration node and clear the contents of the in-memory audit log.

The in-memory portion of the audit log holds approximately 1 MB of audit information. Depending on the command text size and the number of parameters, this equals 1 MB records or approximately 6000 commands.

Once the in-memory audit log reaches maximum capacity, the log is written to a local file on the configuration node in the `/dumps/audit` directory. The **catauditlog** 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 require any special command to decode it.

The in-memory log entries are reset and cleared automatically, ready to accumulate new commands. The on-disk portion of the audit log can then be analyzed at a later date.

The **lsdumps** command with **-prefix** parameter (and the `/dumps/audit` file) can be used to list the files on the disk.

As commands are executed, they are recorded in the in-memory audit log. When the in-memory audit log becomes full, it is automatically dumped to an audit log file and the in-memory audit log is cleared.

An invocation example

This example lists the five most recent audit log entries.

```
catauditlog -delim : -first 5
```

The resulting output:

audit_seq_no	timestamp	cluster_user	challenge	source_panel	target_panel	ssh_ip_address	result	res_obj_id	action_cmd
0	160313152255	superuser		7830619-2	7830619-2		0	0	satask restart
1	160313152303	superuser		01-2	01-1	9.174.187.11	0	0	satask chnode1
2	160313152312	superuser		01-1	01-2	9.174.187.11	0	0	satask chnode1
3	160313152314	superuser		01-1	01-1	9.174.187.11	0	0	satask chnode1
4	160313152316	superuser				9.174.187.11	0	0	svctask chenc1
5	160313152349	superuser				9.174.187.11	0	0	svctask mkmdis
6	160313152352	superuser				9.174.187.11	0	0	svctask mkarra

dumppauditlog

Use the **dumppauditlog** 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

►► — dumppauditlog — ◀◀

Parameters

There are no parameters.

Description

This command dumps the contents of the audit log to a file on the current configuration node in a clustered system (system). It also clears the contents of the audit log. This command is logged as the first entry in the new audit log.

Use this command to manually dump the contents of the in-memory audit log to a file on the current configuration node and clear the contents of the in-memory audit log. Use the **catauditlog** command to display the in-memory 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 system. 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 system. 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 system ID number that does not match the current one are considered to be older than files that match the system 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 clustered system ID at the time that the dump was created
- *source_panel* is the source panel ID in the audit log that is being dumped
- *target_panel* indicates the target panel ID in the audit log that is being dumped

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 **catauditlog** command; however, the **dumpauditlog** command displays the information with one field per line. The **lsdumps** command displays a list of the audit log dumps that are available on the nodes in the clustered system.

A sample audit log entry:

```
Auditlog Entry 0
  Sequence Num      : 0
  Timestamp         : Sun Mar 13 15:22:55 2016
                   : Epoch + 1457882575
  Cluster User      : superuser
  Challenge         :
  SSH IP Address    :
  Result Code       : 0
  Result Obj ID     : 0
  Action Cmd        : satask restartservice -service tomcat
  Source_Panel      : 7830619-2
  Target_Panel      : 7830619-2
```

An invocation example

```
dumpauditlog
```

The resulting output:

No feedback

lsauditlogdumps (Deprecated)

Attention: The **lsauditlogdumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

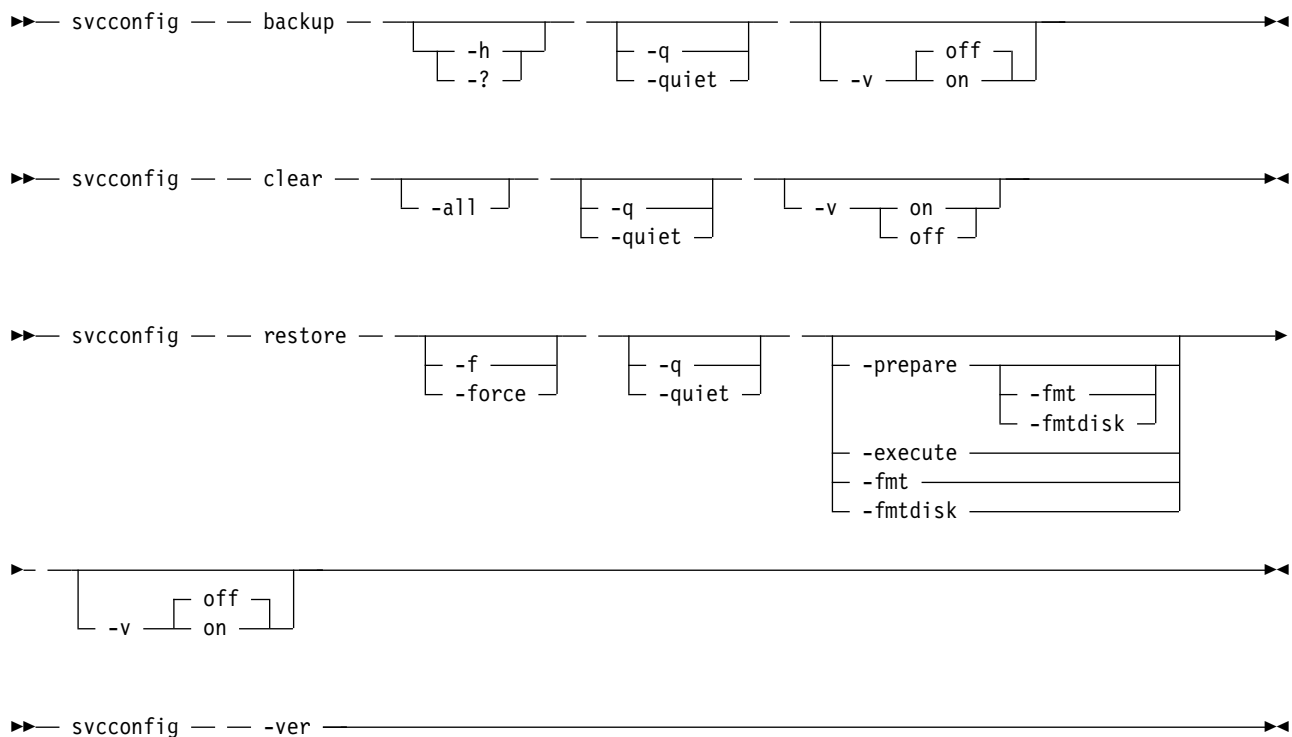
Chapter 5. Backup and restore commands

Backup and restore commands are used for backing up and restoring configuration information with the SAN Volume Controller.

svcconfig

Use the **svcconfig** command help option to obtain summary information about the syntax of the **svcconfig** command and actions. You can enter this command any time after a clustered system (system) is created.

Syntax



Parameters

backup

(Optional) Saves the current clustered system (system) configuration in the /tmp directory.

-quiet

Suppresses standard output (STDOUT) messages from the console.

clear

(Optional) Erases the files in the /tmp directory.

-all

(Optional) Erases all configuration files.

-f | force

(Optional) Forces continued processing where possible.

- q | quiet**
(Optional) Suppresses console output (STDOUT).
- restore**
(Optional) Checks the current configuration against the backup configuration in the /tmp directory.
- prepare -fmt | fmt disk**
(Optional) Verifies the current configuration against the information in svc.config.backup.xml; then prepares commands for processing in svc.config.restore.sh, and then produces a log of events in svc.config.restore.prepare.
- execute**
(Optional) Runs the command script svc.config.restore.sh, and produces a log of events in svc.config.restore.execute.log.
- fmt**
(Optional) Specifies that the volume should be formatted before use. Includes the **-fmt disk** option on all **mkvdisk** commands to be issued. You cannot specify **-fmt** with **-execute**.
- fmt disk**
(Optional) Specifies that the volume should be formatted before use. You cannot specify **-fmt disk** with **-execute**.
- v on | off**
Produces verbose output (on); the default is regular output (off).
- h | -?**
(Optional) Provides general and action-specific help.
- ver**
(Required) Returns the version number for the **svconfig** command.

Description

This command provides syntax help for svconfig.

An invocation example

```
svconfig -ver
svconfig -?
svconfig backup -h
```

backup

Use the **backup** command to back up the configuration. Enter this command any time after creating clustered system (system).

Syntax

```

>>> svconfig -- backup --
                                ┌ -quiet ┐
                                └────────┘
                                ┌ -v ┐ ┌ off ┐
                                └───┘ └ on ┘

```

Parameters

- quiet**
Suppresses standard output (STDOUT) messages from the console.
- v on | off**
Displays normal (off, the default state) or verbose (on) command messages.

Description

The **backup** command extracts and stores configuration information from the system. The **backup** command produces the `svc.config.backup.xml`, `svc.config.backup.sh`, and `svc.config.backup.log` files, and saves them in the `/tmp` folder. 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 the `/tmp` folder, it is archived as `svc.config.backup.bak`; only one archive file is stored in the `/tmp` folder.

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

```
svcconfig backup
```

The resulting output:

No feedback

clear

Use the **clear** command to erase files in the `/tmp` directory that were previously produced by other **svcconfig** commands. You can enter this command any time after a clustered system (system) has been created.

Syntax

```
➤➤ svcconfig — — clear — [ -all ] [ -q | -quiet ] [ -v [ on | off ] ] ➤➤
```

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 configured node.

You can use the **clear** command without the **-all** parameter to erase files of the form:

```
/tmp/svc.config*.sh  
/tmp/svc.config*.log
```

You can use the **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

```
svcconfig clear -all
```

The resulting output:

No feedback

cron

Use the **cron** command to back up the configuration. Enter this command any time after creating the clustered system (system).

Syntax

```
➤— svcconfig — — cron — [ -quiet ] [ -v [ off | on ] ] —————➤
```

Parameters

-q, -quiet

Suppresses standard output (STDOUT) messages from the console.

-v on, -v off

Displays normal (off, the default state) or verbose (on) command messages.

Description

This command generates configuration files and places them in the configuration files directory. The file `svc.config.cron.xml_(node)` contains configuration detail. The file `svc.config.cron.log_(node)` contains a log of events. The file `svc.config.cron.sh_(node)` contains a script of the commands used to determine the configuration.

Any pre-existing file `svc.config.cron.xml_(node)` is archived as `svc.config.cron.bak_(node)`. Only one such archive is kept.

The configuration files directory is `/dumps`.

An invocation example


```
svcconfig cron
svcconfig cron -q
svcconfig cron -v on
```

recover

Use the **recover** command to recover the clustered system configuration in two phases, the preparation phase and the execution phase. This is a component of T3 Recovery.

Syntax

```
➤— svcconfig — — recover — [ -f | -force ] [ -q | -quiet ] [ -prepare | -execute ] —————➤
```



Parameters

-execute

(Optional) Runs the command script `svc.config.recover.sh` and produces a log of events in `svc.config.recover.execute.log`.

-f, -force

(Optional) Forces continued processing where possible.

-prepare

(Optional) Verifies the current configuration against the information in `svc.config.backup.xml` on the configuration to be recovered. Prepares commands for processing in `svc.config.recover.sh`, and produces a log of events in `svc.config.recover.prepare.log`.

-q, -quiet

(Optional) Suppresses console output (STDOUT).

-v on, -v off

(Optional) Produces verbose output (on); the default is regular output (off).

Description

The **recover** command recovers the target system configuration from the `svc.config.backup.xml` file, and associated `.key` files (if present) in the configuration files folder.

The recover operation is performed in two phases: prepare and execute. If neither the **-prepare** nor the **-execute** option is specified, the command performs both phases in sequence, producing only a single event log: `svc.config.recover.log`.

The configuration files directory is `/tmp`.

An invocation example

```


svcconfig recover -prepare
svcconfig recover -execute


```

restore

Use the **restore** command to restore the clustered system (system) to its previous configuration. This command uses the configuration files in the `/tmp` folder .

Syntax





Parameters

- f | force**
(Optional) Forces continued processing where possible.
- q | quiet**
(Optional) Suppresses console output (STDOUT).
- prepare -fmt | fmt disk**
(Optional) Verifies the current configuration against the information in `svc.config.backup.xml`, prepares commands for processing in `svc.config.restore.sh`, and produces a log of events in `svc.config.restore.prepare`.
- execute**
(Optional) Runs the command script `svc.config.restore.sh`, and produces a log of events in `svc.config.restore.execute.log`.
- fmt**
(Optional) Specifies that the volume must be formatted before use. Includes the **-fmt disk** option on all **mkv disk** commands to be issued. You cannot specify **-fmt** with **-execute**.
- fmt disk**
(Optional) Specifies that the volume must be formatted before use. You cannot specify **-fmt disk** with **-execute**.
- v on | off**
(Optional) Produces verbose output (on); the default is regular output (off).

Description

The **restore** command restores the target system configuration from the `svc.config.backup.xml` file in the `/tmp` folder. If neither the **-prepare** nor the **-execute** option is specified, the command performs both phases in sequence, producing only a single event log: `svc.config.restore.log`.

The restore operation is also known as a T4 (Tier 4) Recovery, and can only be used on a system having just been started. The restore operation can not be used on a system having any nonautomatic objects configured, such as storage pools or volumes.

The restore operation is performed in two phases: prepare and execute.

The command pauses for eight minutes if any nodes are added during this process, informing the user of this at run-time.

An invocation example

```
svcconfig restore
```

The resulting output:

No feedback

An invocation example

```
svcconfig restore -prepare -fmt
```

The resulting output:

No feedback

An invocation example

```
svcconfig restore -execute
```

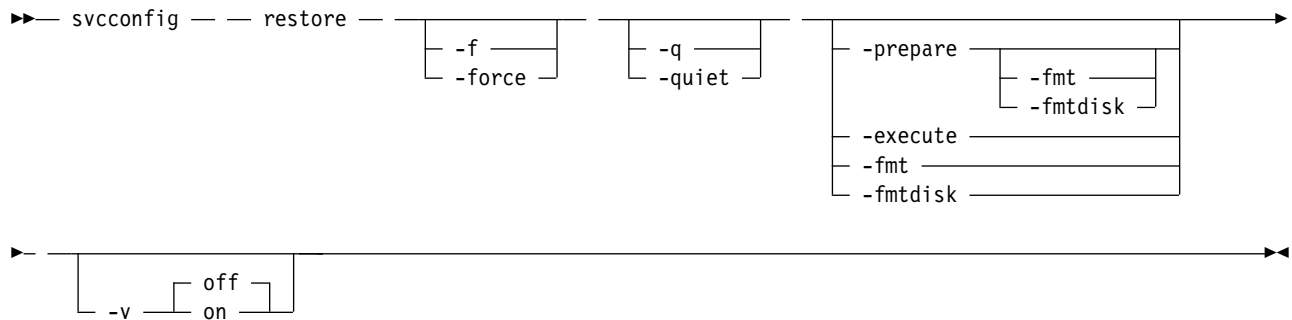
The resulting output:

No feedback

restore

Use the **restore** command to restore the clustered system (system) to its previous configuration. This command uses the configuration files in the /tmp folder .

Syntax



Parameters

- f | force**
(Optional) Forces continued processing where possible.
- q | quiet**
(Optional) Suppresses console output (STDOUT).
- prepare -fmt | fmtdisk**
(Optional) Verifies the current configuration against the information in `svc.config.backup.xml`, prepares commands for processing in `svc.config.restore.sh`, and produces a log of events in `svc.config.restore.prepare`.
- execute**
(Optional) Runs the command script `svc.config.restore.sh`, and produces a log of events in `svc.config.restore.execute.log`.
- fmt**
(Optional) Specifies that the volume must be formatted before use. Includes the **-fmtdisk** option on all **mkvdisk** commands to be issued. You cannot specify **-fmt** with **-execute**.
- fmtdisk**
(Optional) Specifies that the volume must be formatted before use. You cannot specify **-fmtdisk** with **-execute**.
- v on | off**
(Optional) Produces verbose output (on); the default is regular output (off).

Description

The **restore** command restores the target system configuration from the `svc.config.backup.xml` file in the /tmp folder. If neither the **-prepare** nor the **-execute** option is specified, the command performs both phases in sequence, producing only a single event log: `svc.config.restore.log`.

The restore operation is also known as a T4 (Tier 4) Recovery, and can only be used on a system having just been started. The restore operation can not be used on a system having any nonautomatic objects configured, such as storage pools or volumes.

The restore operation is performed in two phases: prepare and execute.

The command pauses for eight minutes if any nodes are added during this process, informing the user of this at run-time.

An invocation example

```
svcconfig restore
```

The resulting output:

No feedback

An invocation example

```
svcconfig restore -prepare -fmt
```

The resulting output:

No feedback

An invocation example

```
svcconfig restore -execute
```

The resulting output:

No feedback

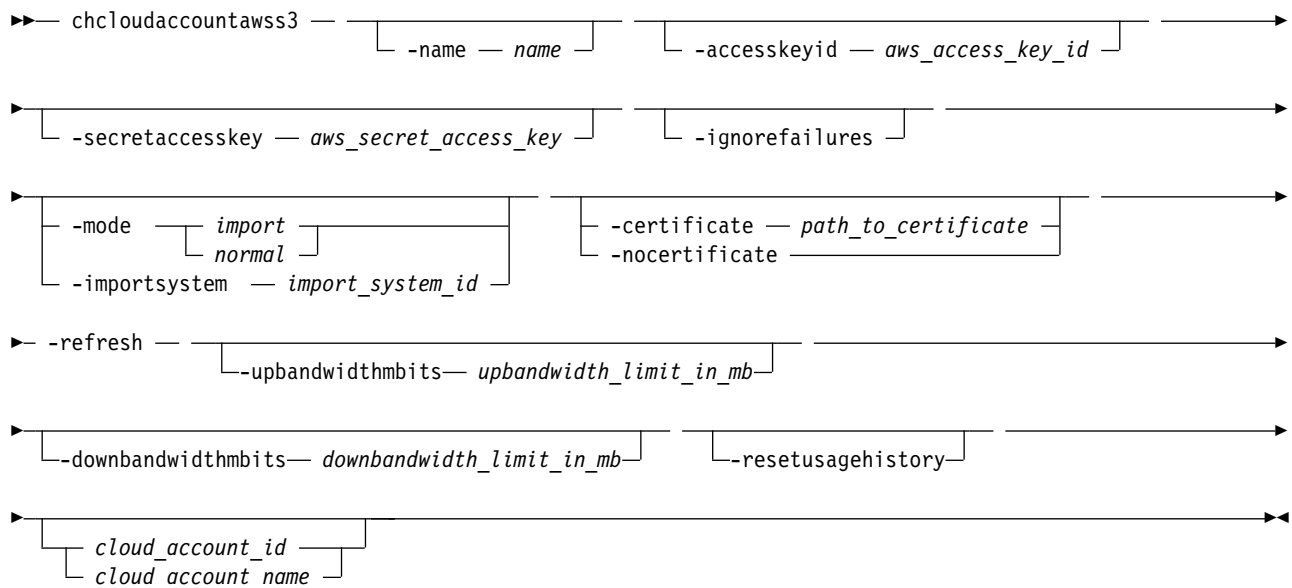
Chapter 6. Cloud account commands

Cloud account commands are used to create, change, or list details about cloud accounts.

chcloudaccountawss3

Use the **chcloudaccountawss3** command to modify the cloud account (that uses Amazon S3 storage) parameters or mode.

Syntax



Parameters

- name** *name*
(Optional) Specifies the new or modified cloud account name. The value must be an alphanumeric value.
- accesskeyid** *aws_access_key_id*
(Optional) Specifies the value for the public part of the Amazon Web Services (AWS) access key. Use this access key to access cloud storage.
- secretaccesskey** *aws_secret_access_key*
(Optional) Specifies the value for the private part of the Amazon Web Services (AWS) access key. This access key is for the AWS user that the system uses to access cloud storage.
- ignorefailures**
(Optional) Changes the access key whether the new access key works.
- mode** *import* | *normal*
(Optional) Specifies the new or modified cloud account mode. The values are *import* or *normal*.
- importsystem** *import_system_id*
(Optional) Specifies that the system's data be imported.

Note: You must specify *-mode import* first.

-certificate *path_to_certificate*

(Optional) Specifies the path for the SSL certificate to use when you authenticate to the new or modified cloud account storage. The value must be an alphanumeric string of 1 - 255 characters (in base64-encoded PEM format).

-nocertificate

(Optional) Specifies that the custom SSL certificate that was used to authenticate to the new or modified cloud account storage be used to stop the system.

-refresh

(Optional) Specifies a refresh of the system import candidates. If the account is in import mode, this parameter specifies a refresh of the data available for import.

-downbandwidthbits *downbandwidth_limit_in_mb*

(Optional) Specifies the download bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

-upbandwidthbits *upbandwidth_limit_in_mb*

(Optional) Specifies the upload bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

-resetusagehistory

(Optional) Resets the usage history (to 0). Storage consumption that reflects the space that is consumed on the cloud account is cumulative, which means that it remains in the current day row (the 0th row).

cloud_account_id | *cloud_account_name*

(Required) Specifies the cloud account ID or name to modify. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command modifies the parameters for the cloud account (created using **mk1cloudaccountawss3**) that uses Amazon S3 storage.

The **-mode** parameter, the **-refresh** parameter, and any of the user credentials parameters groups are mutually exclusive.

This command fails and no changes are made if the supplied credentials do not provide authentication. Credentials include:

- **-accesskeyid**
- **-secretaccesskey**
- **-certificate** or **nocertificate**

For example, if the network is down then the system cannot confirm that a new secret access key is valid, the command fails. Specify **-ignorefailures** to overwrite this feature. If you specify invalid credentials but you do specify **-ignorefailures**, an online account becomes offline and an error is generated in the error log that describes the authentication failure.

The secret access key is sensitive system information and is stored in encrypted form. It is not available in the system dumps and in the audit log it is replaced with six hash ("#") symbols.

If you specify this command against an offline account and these new details enable the account to start working (for example, you enter an expired password) the account becomes online.

You can only change the mode if the account is not being used by any system volumes. A mode change requires the account to be online and the system be able to communicate with the cloud server.

Note: You can have a maximum of:

- One cloud account per clustered system (system)
- 1024 volumes with cloud snapshots enabled
- 256 cloud snapshots per volume
- 512 volume groups

An invocation example

```
chcloudaccountawss3 -name myamazon cloudaccount0
```

The resulting output:

No feedback

An invocation example

```
chcloudaccountawss3 -mode import -importsystem 000002007D40A162 0
```

The resulting output:

No feedback

An invocation example

```
chcloudaccountawss3 -upbandwidthmbits 100 -downbandwidthmbits 100 cloudaccount0
```

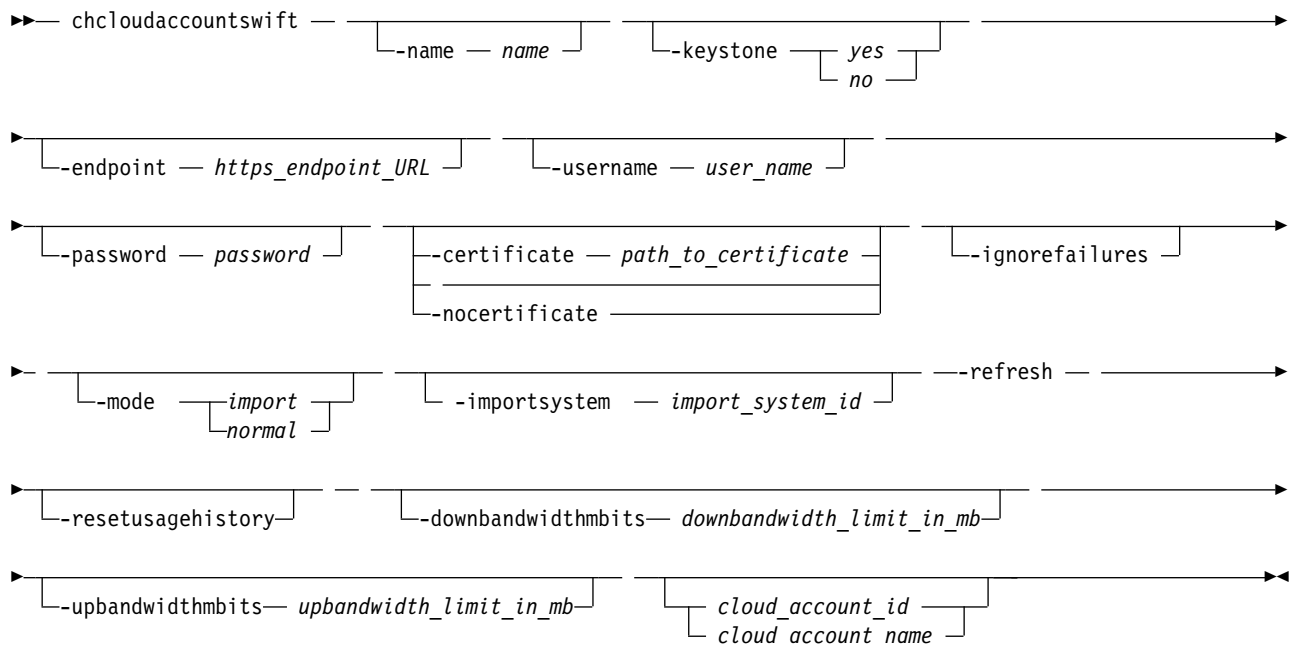
The resulting output:

No feedback

chcloudaccountswift

Use the **chcloudaccountswift** command to modify the cloud account (that uses OpenStack Swift storage) parameters or mode.

Syntax



Parameters

- name** *name*
(Optional) Specifies the new or modified OpenStack name you must use to access cloud account storage. The value must be an alphanumeric value.
- keystone** *yes* | *no*
(Optional) Specifies that keystone authentication be used. The values are *yes* or *no*.
- endpoint** *https_endpoint_URL*
(Optional) Specifies the URL (that the system uses to access object storage) to change for the cloud account. If OpenStack Keystone authentication is used, the URL specified must be the URL for the Keystone authentication. If Keystone authentication is not used, the URL specified must be the URL for the Swift account. The value must be 8 - 128 characters and must be a valid URL address.
- username** *user_name*
(Optional) Specifies the OpenStack user name that the system must use to access cloud account storage.
- password** *password*
(Optional) Specifies the password value to use to authenticate to cloud storage. For IBM SoftLayer[®] accounts, this password is the application programming interface (API) key. The value must be 1 - 64 alphanumeric characters and it must not begin or end with a space.
- certificate** *path_to_certificate*
(Optional) Specifies the path for the SSL certificate to use when you authenticate to new or modified cloud account storage. The value must be an alphanumeric string of 1 - 255 characters (in base64-encoded PEM format).
- nocertificate**
(Optional) Specifies that the custom SSL certificate that was used to authenticate to the new or modified cloud account storage be used to stop the system.
- ignorefailures**
(Optional) Specifies that the access key be changed whether the new access key works.
- mode** *import* | *normal*
(Optional) Specifies the new or modified cloud account mode. The values are *import* or *normal*.
- importsystem** *import_system_id*
(Optional) Specifies that the system's data be imported.

Note: You must specify *-mode import* first.
- refresh**
(Optional) Specifies a refresh of the system import candidates. If the account is in *import* mode, this specifies a refresh of the data available for import.
- downbandwidthmbits** *downbandwidth_limit_in_mb*
(Optional) Specifies the download bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.
- upbandwidthmbits** *upbandwidth_limit_in_mb*
(Optional) Specifies the upload bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.
- resetusagehistory**
(Optional) Resets the usage history (to 0). Storage consumption that reflects the space that is consumed on the cloud account is cumulative, which means that it remains in the current day row (the 0th row).

cloud_account_id | *cloud_account_name*

(Required) Specifies the cloud account ID or name to modify. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command modifies the parameters for the cloud account (created by using **mkcloudaccountswift**) that uses OpenStack Swift storage.

At least one parameter must be set.

The **-mode** parameter, the **-refresh** parameter, and any of the user credentials parameters groups are mutually exclusive. Credentials include:

- **-keystone**
- **-endpoint**
- **-username**
- **-password**
- **-certificate** or **nocertificate**

The command fails if the supplied authentication credentials are unsuccessful. For example, if the network is down the system cannot confirm that the **secretaccesskey** is valid (and the command fails). Specify **-ignorefailures** to override this. If you specify incorrect credentials and the **-ignorefailures** parameter, an online account becomes offline and an error is generated in the log describing the authentication failure.

The password is treated as sensitive system information. It is stored in an encrypted form and not available in system dumps. In the audit log it is replaced with six hash ("#") symbols.

If a certificate is supplied and the command succeeds the certificate file is deleted from the local file system.

If you specify this command against an offline account and these new details enable the account to start working (for example, you enter an expired password) the account becomes online.

You can change the mode if the account is not being used by any system volumes. A mode change requires the account to be online and the system be able to communicate with the cloud server.

Note: You can have a maximum of:

- One cloud account per clustered system (system)
- 1024 volumes with cloud snapshots enabled
- 256 cloud snapshots per volume
- 512 volume groups

An invocation example

```
chcloudaccountswift -certificate /tmp/new-cert.pem -ignorefailures mysift
```

The resulting output:

No feedback

An invocation example

```
chcloudaccountswift -mode import -importsystem 000002007D40A162 0
```

The resulting output:

No feedback

An invocation example

```
chcloudaccountawss3 -username newuser -password simpsons 0
```

The resulting output:

No feedback

An invocation example

```
chcloudaccountswift -upbandwidthmbits 100 -downbandwidthmbits 100 cloudaccount0
```

The resulting output:

No feedback

Iscloudaccount

Use the **lscloudaccount** command to display information about the configured cloud accounts.

Syntax

```
➔ lscloudaccount — [ -nohdr ] [ -delim delimiter ] [ cloud_account_id | cloud_account_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 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 1-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.

cloud_account_id | *cloud_account_name*

(Optional) Specifies the name or ID for the detailed view of the account.

Description

This command displays information about configured cloud accounts.

This table provides the attribute values that can be displayed as output view data.

Table 23. *Iscloudaccount* output

Attribute	Description
id	Indicates the cloud account ID. The value is a number.

Table 23. *Iscloudaccount output (continued)*

Attribute	Description
name	Indicates the cloud account name. The value is an alphanumeric string.
type	Indicates the cloud account provider. The values are awss3 or swift.
status	Indicates the cloud account status. The values are online or offline.
mode	Indicates the cloud account mode. The values are normal or import.
active_volume_count	Indicates the number of volumes in the system that use the account. The value must be a number.
backup_volume_count	Indicates the number of volumes that are backed up to the cloud account. The value must be a number.
import_system_id	Indicates the system ID for the system from where the data is being imported. The value must be a 16-character uppercase hexadecimal number (or blank).
import_system_name	Indicates the system name from where the data is being imported. The value must be an alphanumeric string (or blank).
error_sequence_number	Indicates an error (for offline accounts). The value must be a number (or blank).
refreshing	Indicates whether the system is refreshing its cloud storage view (for import mode accounts). The values are yes or no.
backup_timestamp	Indicates the timestamp for the most recent backup. The value must be in the format YYMMDDHHMMSS (or blank).
certificate	Indicates whether SSL is configured for an account that uses certificates. The values are yes or no.
certificate_expiry	Indicates the time and date that a certificate expires. The value must be blank or be in this format: Dec 7 10:07:59 2015 GMT
endpoint	Indicates the endpoint URL for swift accounts. The value must be a valid URL (or blank).
awss3_bucket_prefix	Indicates the bucket prefix that is being used for S3 accounts. The value must be a valid bucket prefix (or blank).
awss3_access_key_id	Indicates the user access key ID for S3 accounts. The value must be a valid access key ID (or blank).
awss3_region	Indicates the region chosen for cloud storage for S3 accounts. The value must be for a valid AWS region (or blank).
swift_keystone	Indicates whether keystone authentication is in use. The value must be yes or no.
swift_container_prefix	Indicates the container prefix for Swift accounts. The value must be a valid container prefix or blank.
swift_tenant_name	Indicates the tenant name that is used for authentication for swift accounts. The value must be a valid tenant name (or blank).
swift_user_name	Indicates the user name that is used for authentication for swift accounts. The value must be a valid user name (or blank).
encrypt	Indicates the encryption status for the cloud account. The values are yes and no.

A concise invocation example

```
lscloudaccount
```

The resulting output:

```
id name      type status mode  active_volume_count backup_volume_count import_system_id import_system_name error_seque
0 importer  swift online import 2                123                000002007D40A162 cluster1
```

A detailed invocation example

```
lscloudaccount 1
```

The resulting output:

```
id 0
name varyja
type swift
status online
mode normal
active_volume_count 0
backup_volume_count 1
import_system_id
import_system_name
error_sequence_number
refreshing no
backup_timestamp 151021114002
certificate yes
certificate_expiry Dec 7 10:07:59 2015 GMT
endpoint https://thesecurecloud.company.com:4000/auth/v3.0
awss3_bucket_prefix
awss3_access_key_id
awss3_region
swift_keystone yes
swift_container_prefix svc-1
swift_tenant_name mytenant
swift_user_name storeman
```

lscloudaccountusage

Use the **lscloudaccountusage** command to list usage information about configured cloud storage accounts.

Syntax

```
➤ lscloudaccountusage - [ -nohdr ] - [ -delim delimiter ]
➤ [ cloud_account_id ]
  [ cloud_account_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 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 1-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.

`cloud_account_id` | `cloud_account_name`

(Optional) Specifies the cloud account to list details for. The ID value must be a number and the name value must be an alphanumeric string.

Description

This command displays usage information about configured cloud storage accounts. The information involves chargeable resource usage.

This table provides the attribute values that can be displayed as output view data.

Table 24. *lscloudaccountusage* output

Attribute	Description
id	Indicates the ID for the cloud account. The value must be a number 0 - 4294967295.
name	Indicates the name for the cloud account. The value must be an alphanumeric string.
date	Indicates the date for the system data that is displayed. Each row shows usage of one day. The value must be in YYYYMMDD format. This value is computed relative to the current configured system date. The date in very first entry equals the current date. If you manually change the system date, the changes are not reflected in the output for the date field until midnight (the time 00:00). If you change your system date got to accommodate time zone changes, it is reflected in the output instantly. After midnight, any subsequent entry is for the next 24 hour period.
upload_data_mb	Indicates the uploaded data for one day. The value must be a number 0 - 18446744073709551615.
download_data_mb	Indicates the downloaded data for one day. The value must be a number 0 - 18446744073709551615.
storage_consumed_gb	Indicates the volume of data that is stored in this cloud account. The value must be a number 0 - 18446744073709551615.

Note: For a detailed view, there are 180 rows. Each row has information corresponding to one full day, and every field reflects activity for that day except for the `storage_consumed_gb`, which is cumulative. The most recent entry reflects the current day.

An invocation example

```
lscloudaccountusage
```

The resulting output:

id	name	date	upload_data_mb	download_data_mb	storage_consumed_gb
0	cloudaccount0	20151023	194560	900	6700
1	cloudaccount1	20151023	204800	1500	10700

An invocation example

```
lscloudaccountusage 0
```

The resulting output:

id	name	date	upload_data_mb	download_data_mb	storage_consumed_gb
0	cloudaccount0	20151023	194560	900	6687
0	cloudaccount0	20151022	3584000	150	6495
0	cloudaccount0	20151021	1024	17152	3010

Iscloudaccountimportcandidate

Use the **iscloudaccountimportcandidate** command to list information about systems that have data that is stored in the cloud accounts that are defined on this system.

Syntax

```
➤— iscloudaccountimportcandidate —┬─ -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 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 1-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 lists information about systems that have data that is stored in the cloud accounts that are defined on this system.

This command provides information about the valid options for **chcloudaccount -import**. To refresh the view by reloading what is on the cloud server, specify **chcloudaccount -refresh**.

This table provides the attribute values that can be displayed as output view data.

Table 25. Iscloudaccountimportcandidate output

Attribute	Description
cloud_account_id	Indicates the ID for the cloud account that contains data from another system. The value must be a numeric string.
cloud_account_name	Indicates the name for the cloud account that contains data from another system. The value must be an alphanumeric string.
import_system_id	Indicates the system ID of the system that has data on the cloud account. The value must be a 16-character string in hexadecimal uppercase.
import_system_name	Indicates the system name of the system that has data on the cloud account. The value must be an alphanumeric string.
backup_volume_count	Indicates the number of volumes that are backed up by the imported system. The value must be a numeric string.
backup_size	Indicates the approximate amount of cloud storage that is in use by snapshots from the imported system.
backup_timestamp	Indicates the timestamp of the most recent volume backup (by the other system). The value must be in the YYMMDDHHMMSS format or be blank. This is displayed in UNIX time.

An invocation example

lscloudaccountimportcandidate

The resulting output:

cloud_account_id	cloud_account_name	import_system_id	import_system_name	backup_volume_count	backup_size	backup_timestamp
0	my_amazon	00002007D40A162	cluster1	0	0.00GB	
0	my_amazon	00002007F42E813	cluster2	44	15.25TB	151008084203

mkcloudaccountawss3

Use the **mkcloudaccountawss3** command to configure a new cloud account that uses Amazon S3 object storage.

Syntax

```
➤➤ mkcloudaccountawss3 — -name name — -bucketprefix bucket_prefix — ➤➤
➤ -accesskeyid aws_access_key_id — -secretaccesskey aws_secret_access_key — ➤➤
➤ -certificate path_to_certificate -upbandwidthmbits upbandwidth_limit_in_mb — ➤➤
➤ -downbandwidthmbits downbandwidth_limit_in_mb -region aws_region — ➤➤
➤ -encrypt yes no — ➤➤
```

Parameters

- name** *name*
(Optional) Specifies the name for the cloud account. The value must be an alphanumeric string.
- bucketprefix** *bucket_prefix*
(Required) Specifies the prefix for the S3 bucket names that the system uses. The value must be a lower-case alphabetic string 3 - 58 characters long (with no dot or period at the end of the string, and no dot or period next to another dot or period).
- accesskeyid** *aws_access_key_id*
(Required) Specifies the public part of the Amazon Web Services (AWS) access key credential of the AWS user that the system use to access the cloud storage. The value must be a 20-character alphanumeric string of uppercase letters and numbers.
- secretaccesskey** *aws_secret_access_key*
(Required) Specifies the non-public part of the AWS access key credential that the system use to access the cloud storage. The value must be a 40-character alphanumeric string (That can contain slashes, or "/").
- certificate** *path_to_certificate*
(Optional) Specifies the path to an SSL certificate authority (CA) certificate for AWS S3. The value must be an alphanumeric string of 1 - 255 characters (in base64-encoded PEM format).
- upbandwidthmbits** *upbandwidth_limit_in_mb*
(Optional) Specifies the upload bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

-downbandwidthmbits *downbandwidth_limit_in_mb*

(Optional) Specifies the download bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

-region *aws_region*

(Optional) Specifies the AWS region to use to access the cloud account and store data.

-encrypt *yes | no*

(Optional) Specifies whether to encrypt the data in the cloud account. By default, encryption is enabled unless you specify **-encrypt no**.

Description

This command configures a new cloud account that uses Amazon S3 object storage.

An invocation example

```
mkcloudaccountaws3 -name myamazon
                   -bucketprefix svc_backups
                   -accesskeyid AKIAIOSFODNN7EXAMPLE
                   -secretaccesskey wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
                   -upbandwidthmbits 100
                   -downbandwidthmbits 100
```

The resulting output:

Cloud Account, id [0], successfully created

Note: If a cluster contains an encrypted cloud account that uses USB encryption, a USB flash drive with the cluster master key must be present in the configuration node before the cloud account can move to the online state. This requirement is necessary when the cluster is powered down, and then restarted.

mkcloudaccountswift

Use the **mkcloudaccountswift** command to configure a new cloud account that uses OpenStack Swift object storage (including IBM SoftLayer).

Syntax

```
➤ mkcloudaccountswift — name — keystone —————➤
➤ --containerprefix container_prefix — --endpoint http_endpoint_URL —————➤
                                     https_endpoint_URL
➤ --tenantname tenant_name — --username user_name — --password password —————➤
➤ certificate path_to_certificate — upbandwidthmbits upbandwidth_limit_in_mb —————➤
➤ downbandwidthmbits downbandwidth_limit_in_mb — encrypt yes —————➤
                                                         no
```

Parameters

-name *name*

(Optional) Specifies the account identifier. The value must be an alphanumeric string.

-keystone

(Optional) Specifies that the system authenticates with OpenStack Keystone. If you do not specify this parameter, the system authenticates with OpenStack TempAuth.

-containerprefix *container_prefix*

(Required) Specifies the Swift container names the system uses or creates. The value must be 1 - 12 characters and contain no spaces or slashes.

-endpoint *http_endpoint_URL* | *https_endpoint_URL*

(Required) Specifies the URL that the system uses to access object storage.

If Keystone authentication is used, this should be the URL of the Keystone service, probably ending with v2.0. Otherwise this should be the URL of the Swift service.

-tenantname *tenant_name*

(Required) Specifies the OpenStack tenant the system uses to access cloud storage. The value must be a 1 - 64 alphanumeric characters that contain no spaces.

-username *user_name*

(Required) Specifies the OpenStack user name the system uses to access cloud storage. The value must be 1 - 255 alphanumeric characters with no spaces.

-password *password*

(Required) Specifies the password that the system uses to access cloud storage. For IBM SoftLayer accounts, this password is the application programming interface (API) key. The value must be 1 - 64 alphanumeric characters and it must not begin or end with a space.

-certificate *path_to_certificate*

(Optional) Specifies the file path for the object storage server SSL certificate. The value must be:

- 1 - 255 alphanumeric characters with no period or dot next to another period or dot, and no period or dot at the start or end of the specified value
- In base64-encoded PEM format

-upbandwidthmbits *upbandwidth_limit_in_mb*

(Optional) Specifies the upload bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

-downbandwidthmbits *downbandwidth_limit_in_mb*

(Optional) Specifies the download bandwidth limit in megabits per second (Mbps). The value must be a number 1 - 10240.

-encrypt *yes* | *no*

(Optional) Specifies whether to encrypt the data in the cloud account. By default, encryption is enabled unless you specify **-encrypt no**.

Description

This command configures a new cloud account that uses OpenStack Swift object storage.

An invocation example

```
mkcloudaccountswift -containerprefix svc_backups
                    -endpoint https://lon02.objectstorage.softlayer.net/auth/v1.0
                    -tenantname mytenant
                    -username jamivard
                    -password WKf84FAQRKLOICDF53LANBWKf84FAQRKLOICDF53LANBEXAMPLEEXAMPLEEXAMPL
                    -upbandwidthmbits 100
                    -downbandwidthmbits 100
```

The resulting output:

```
Cloud Account, id [0], successfully created
```

Note: If a cluster contains an encrypted cloud account that uses USB encryption, a USB flash drive with the cluster master key must be present in the configuration node before the cloud account can move to the online state. This requirement is necessary when the cluster is powered down, and then restarted.

rmcloudaccount

Use the **rmcloudaccount** command to delete a cloud account from the system.

Syntax

```
►►— rmcloudaccount — —┐ cloud_account_id —————►◄◄
                       └ cloud_account_name ┘
```

Parameters

cloud_account_id | *cloud_account_name*

(Required) Specifies the cloud account to remove. The ID value must be a number and the name value must be an alphanumeric string.

Description

This command deletes a cloud account from the system. If no systems have volume data that is stored in the account, the containers are deleted from the cloud storage.

As long as there are no volumes on this system that are using the cloud account, the command deletes the account. If there are no volumes that are left in the account, the system attempts to delete its containers. If there is a problem connecting to the cloud server, the containers are not deleted. If the command times out, the deletion proceeds asynchronously and the account object is removed.

An invocation example

```
rmcloudaccount VardyAmazAcct
```

The resulting output:

No feedback

testcloudaccount

Use the **testcloudaccount** command to run diagnostics against the cloud account and report status on the results.

Syntax

```
►►— testcloudaccount — —┐ cloud_account_id —————►◄◄
                       └ cloud_account_name ┘
```

Parameters

cloud_account_id | *cloud_account_name*

(Required) Specifies the cloud account to test. The ID value must be a number and the name value must be an alphanumeric string.

Description

This command runs diagnostics against the cloud account and reports status. This includes network connectivity, authentication, and cloud storage usage.

This command can be run against an online or offline account.

- If the command is run successfully against an offline account, the account becomes online.
- If the command is run unsuccessfully against an online account, the account becomes offline.

An invocation example

```
testcloudaccount MyVardyAccount
```

The resulting output:

```
Cloud Account, id [0], successfully tested
```


Chapter 7. Clustered system commands

Clustered system (system) commands are used to monitor and modify systems.

addnode (SAN Volume Controller only)

Use the **addnode** command to add a new (candidate) node to an existing clustered system (system). Enter this command any time after a system has been created. If you are adding a new node to a system, make sure the model type of the new node is supported by the SAN Volume Controller system code (code) version of the system. If the model type is not supported by the code, upgrade the system to a code version that supports the model type of the new node.

Syntax



Parameters

-panelname *panel_name*

(Required if you do not specify the **-wwnodename** parameter) Specifies the node that you want to add to a system by the name that is displayed on the display panel. You cannot use this parameter with the **-wwnodename** parameter.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

-wwnodename *wwnn_arg*

(Required if you do not specify the **-panelname** parameter) Specifies the node that you want to add to the system 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 system. 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 system 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 system 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.

-site *site_name* | *site_id*

(Optional) Specifies the numeric site value or site name of the new node.

If the system topology is hyperswap and the I/O group has a configured node, this new node must be located within the same site. If there are no configured nodes in the I/O group (but there are volumes defined in the I/O group that are in active-active relationships) this new node must be located within the same site as any node that was previously in that I/O group

Remember:

- This parameter must be specified whether the system topology is set to stretched or hyperswap.
- If the system topology is stretched and the I/O group has a configured node, this new node must be in another site location.

Description

Note: The **addnode** command is a SAN Volume Controller command. For Storwize V7000, use the **addcontrolenclosure** command.

This command adds a node to the system. You can obtain a list of candidate nodes (nodes that are not already assigned to a system) by typing **lsnodecandidate**. You cannot add a node with less memory than any potential partner nodes that are in the I/O group.

Note: The **lsnodecandidate** command is a SAN Volume Controller command. For Storwize V7000, use the **lscontrolenclosurecandidate** command.

You cannot use this command if the new node is:

- Not capable of encryption but the existing I/O group partner is capable of encryption
- Not capable of encryption but there are storage pools with encryption keys that include MDisks that are not self-encrypting
- Capable of encryption but the node has no encryption license

Note: This command is successful only if the node-enclosure system ID matches the system, or is blank.

Before you add a node to the system, 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 system.

- Is the new node being used to replace a failed node in the system?
- Does the node being added to the system use physical node hardware that has been used as a node in another system, and are both system 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 **lsnode** or the management GUI to determine the WWNN of the system nodes.
2. Shut down all of the hosts that use the system, before you add the node back into the system.
3. Add the node back to the system 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 system, you can do the following:
 - a. On all of the hosts that are connected to the system, unconfigure the Fibre Channel adapter device driver, the disk device driver, and the multipathing driver before you add the node to the system.
 - b. Add the node to the system, 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 system, take the following actions:

1. Ensure that the model type of the new node is supported by the SAN Volume Controller of code for the system. If the model type is not supported by the system code, you must upgrade the system to a version of code that supports the model type of the new node.
2. Record the node serial number, the WWNN, all WWPNNs, 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 clustered system.

Note: Transparent cloud tiering can be enabled on a system if every node on the system supports it. If a system supports transparent cloud tiering, you cannot add nodes that do not support it to the system.

Other considerations when you add a node to a system:

When you add a node to the system using the **addnode** command or the system GUI, you must confirm whether the node has previously been a member of the system. 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 system using the **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 system without corrupting the data.

When a node is added to a system, it displays a state of adding. It can take as long as 30 minutes for the node to be added to the system, particularly if the version of code 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, or if the system has a configuration that exceeds the limits for the node being added. Specify a different name for the node being added.

An invocation example

```
addnode -wwnodename 5005076801e08b -iogrp io_grp0
```

The resulting output:

```
Node, id [6], successfully added
```

An invocation example

```
addnode -panelname 123456 -iogrp 1 -site 2
```

The resulting output:

```
Node, id [6], successfully added
```

An invocation example

```
addnode -wwnodename 5005076801e08b -iogrp io_grp0 -site site1
```


The resulting output:

```
Node, id [6], successfully added
```

addiscsistorageport

Use the **addiscsistorageport** command to establish Internet Small Computer Systems Interface (iSCSI) login sessions from any (or all) nodes in a specified I/O group (or the entire clustered system) to a discovered backend target Internet Small Computer Systems Interface (iSCSI) controller.

Syntax



Parameters

-iogrp *iogrp_id* | *iogrp_name*

(Optional) Specifies I/O group ID or name that is added. The *iogrp_id* value must be 0, 1, 2, or 3. The *iogrp_name* value must be an alphanumeric string.

Specifying this parameter triggers discovery through both nodes for the specified I/O group. The port number on each node (used to establish a session) is displayed in the selected row of discovery results from specifying **detectiscsistorageportcandidate**.

-username *target_user_name*

(Optional) Specifies the target controller user name that is added. The value must be an alphanumeric string up to 256 characters.

If the target controller requires a *target_user_name* and *target_chap* for discovery, the user name for the target controller must be specified.

Some controllers might require that you use the iSCSI qualified name (IQN) user name for discovery. Each nodes IQN is picked up automatically and used if required.

-chapsecret *target_chap*

(Optional) Specifies the *target_chap* required for discovery of the target iSCSI controller that is being added. The value must be an alphanumeric string (case-sensitive) up to 79 characters. This keyword is required if **-username** is specified.

-site *site_id* | *site_name*

(Optional) Specifies the site ID or site name of the host that is being detected. The site ID must be 1 (the default) or 2. The site name must be an alphanumeric value.

Important: This parameter must be specified for a HyperSwap or stretched system.

For a HyperSwap or stretched cluster topology, the site ID must be specified to make sure the session that the attempts to establish a session are made from nodes in the same site as the storage controller.

candidate_id

(Required) Indicates the row ID indicating the selected row in the **lsiscsistorageportcandidate** output.

Description

This command establishes iSCSI login sessions from a specified I/O group (or, if I/O group is not specified, the entire clustered system) to a discovered backend iSCSI target controller.

To use this command, you must first:

1. Specify **detectiscsistorageportcandidate** to detect or discover the backend controller target ports.
2. Specify **lsiscsistorageportcandidate** to display discovery output and locate a unique iSCSI Qualified Name (IQN) and Internet Protocol (IP) combination on the discovered backend storage controller.
3. Specify **addiscsistorageport**. You can establish sessions (from all nodes in a single I/O group) to the backend controller port by specifying the I/O group number of the source port that the iSCSI sessions are initiated from.

If you do not specify an I/O group, sessions are established from all nodes in the system. The source port identifier is in the discovery results. Because some iSCSI controllers represent logical unit numbers (LUNs) as IQNs and might require a different *target_user_name* and *target_chap* for every IQN, these values can also be specified for authentication when establishing the session.

A detailed invocation example

This example shows target discovery that uses an IPv4 IP address for a target iSCSI controller through source port ID 0. For example, first you might specify:

```
detectiscsistorageportcandidate -targetip 192.168.81.91 -srcportid 3 -chapsecret Vardy -site 1
```

You would then specify **lsiscsistorageportcandidate** to list iSCSI port information:

id	src_port_id	target_ip	target_ip	target_iscsiname	io	group_list	configured	status
0	4	192.168.213.33		IQN1	1:1:1:1		yes	full

Then establish a session using **addiscsistorageport** for discovery output row 0.

Note: If you specify **-username** or **-chapsecret** with **detectiscsistorageportcandidate** during discovery, you must specify **-username** or **-chapsecret** for session establishment.

```
addiscsistorageport 0
```

The detailed resulting output:

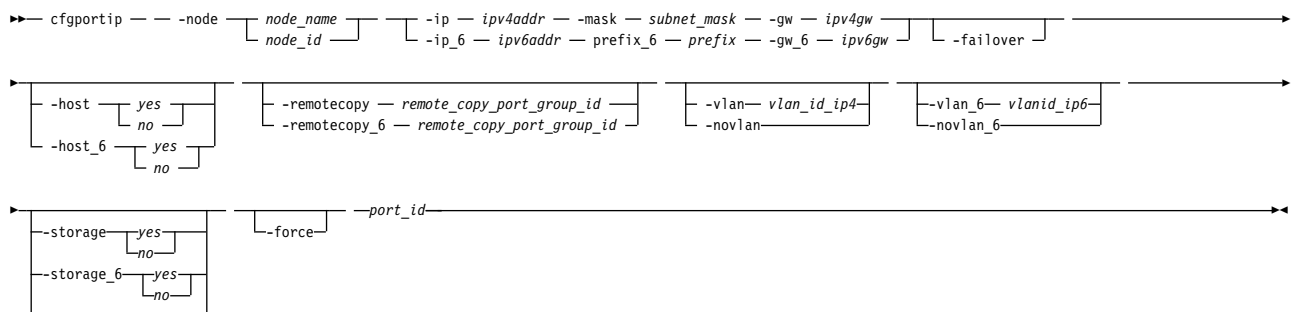
No feedback

cfgportip


Use the **cfgportip** command to assign an Internet Protocol (IP) address to each node Ethernet port for Internet Small Computer System Interface (iSCSI) input/output (I/O).

Syntax

For Internet Protocol Version 4 (IPv4) and Internet Protocol Version 6 (IPv6):



For maximum transmission unit (MTU):



Parameters

-node *node_name* | *node_id*

(Required) Specifies which node has the Ethernet port that the IP address is being assigned to.

Note: This parameter is required for setting a port IP address. It cannot be used with the **-mtu** parameter.

-ip *ipv4addr*

(Required) Sets the Internet Protocol Version 4 (IPv4) address for the Ethernet port. You cannot use this parameter with the **ip_6** parameter.

-ip_6 *ipv6addr*

(Required) Sets the Internet Protocol Version 6 (IPv6) address for the Ethernet port. You cannot use this parameter with the **ip** parameter.

-gw *ipv4addr*

(Required) Sets the IPv4 gateway IP address. You cannot use this parameter with the **gw_6** parameter.

-gw_6 *ipv6gw*

(Required) Sets the IPv6 default gateway address for the port. You cannot use this parameter with the **gw** parameter.

-mask *subnet_mask*

(Required) Sets the IPv4 subnet mask. You cannot use this parameter with the **prefix_6** parameter.

-prefix_6 *prefix*

(Required) 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 not configured or 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.

If the partner node is online, do not use this option.

-mtu *mtu* | **-defaultmtu**

(Required) Specifies the maximum transmission unit (MTU). The default is 1500, with a maximum of 9000. An MTU of 9000 lets you to save CPU utilization for packets of 4 KB and over in size. The increased MTU provides you with improved Internet Small Computer System Interface (iSCSI) performance. Specify **-defaultmtu** to use the default value.

Note: This parameter:

- Must be used when you are setting the clustered system (system) MTU This parameter
- Cannot be used with the **-node** parameter

-iogrp *iogrp*

(Optional) Specifies the I/O group that contains the nodes to modify.

-host *yes* | *no*

(Optional) Specifies the IPv4 address that is used for host attach (the existing system settings are retained). Specifying:

- *yes* reports the IPv4 address to hosts during target discovery (default)
- *no* turns off this report (IPv4 addresses are not reported during host discovery).

-remotecopy *remote_copy_port_group_id*

(Optional) Specifies the IPv4 address that is used for the remote copy function. Remote copy includes

Metro Mirror, Global Mirror, and HyperSwap. It also specifies the ID for the associated port group. These IDs are numerical values (0, 1, or 2) that specify that IP addresses on a system can be part of a partnership for a login. To form a login, IP addresses must be in the same port group. The default is 0, which indicates that the port is not available for partnerships.

Important: To add or delete ports to or from a replication group, make sure the partnership that uses that port group is in a stopped state.

-host_6 *yes | no*

(Optional) Specifies the IPv6 address that is used for host attach (the existing system settings are retained). Specifying:

- *yes* reports the IPv6 address to hosts during target discovery (default).
- *no* turns off this report (IPv6 addresses are not reported during host discovery).

Note: Turning off host attach settings for an IP address that is set to *yes* is disruptive because all host iSCSI sessions to that IP address are logged out.

-remotecopy_6 *remote_copy_port_group_id*

(Optional) Specifies the IPv6 address that is used for the remote copy function. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap. It also specifies the ID for the associated port group. These IDs are numerical values (0, 1, or 2) that specify that IP addresses on a system can be part of a partnership for a login. To form a login, IP addresses must be in the same port group. The default is 0, which indicates that the port is not available for partnerships.

Important: To add or delete ports to or from a replication group, make sure the partnership that uses that port group is in a stopped state.

-vlan *vlanid_ip4*

(Optional) Sets the virtual local area network (VLAN) ID for a IPv4 address that is configured for iSCSI host attach or remote copy function. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap. The VLAN ID for an IPv4 type address can be specified only if the IP address for that port is set. VLAN tagging is disabled for any IP address, so a VLAN ID must be specified by using **-vlan** to turn on VLAN tagging.

Remember: Use **-vlan** with caution. You can:

- Reset VLAN settings, which can disrupt port communication (connection) with hosts or systems (including resetting the VLAN ID for an active iSCSI or IP partnership)
- Reset a VLAN value for a port that does not have VLAN tagging or does not have a configured IP address

The VLAN ID can be set for the failover port that uses the **-failover** attribute.

-novlan

(Optional) Disables VLAN tagging for an IPv4 address for an Ethernet port (which means no VLAN tag is associated with that port).

Remember: Use **-novlan** with caution:

- Resetting VLAN settings can disrupt port communication (connection) with hosts or systems, including resetting the VLAN ID for an active iSCSI or IP partnership.
- You can reset a VLAN value for a port that does not have VLAN tagging or does not have a configured IP address.

-vlan_6 *vlanid_ip6*

(Optional) Sets the virtual local area network (VLAN) ID for a IPv4 address that is configured for iSCSI host attach or remote copy function. Remote copy includes Metro Mirror, Global Mirror, and

HyperSwap. The VLAN ID for an IPv6 type address can be specified only if the IP address for that port is set. VLAN tagging is disabled for any IP address, so a VLAN ID must be specified by using **-vlan** to turn on VLAN tagging.

Remember: Use **-vlan** with caution:

- Resetting VLAN settings can disrupt port communication (connection) with hosts or systems, including resetting the VLAN ID for an active iSCSI or IP partnership.
- You can reset a VLAN tag for a port that does not have VLAN tagging or does not have a configured IP address.

The VLAN ID can be set for the failover port by using the **-failover** attribute.

-novlan_6

(Optional) Disables Virtual local area network (VLAN) tagging for an IPv6 address for an Ethernet port (which means no VLAN tag is associated with that port).

Remember: Use **-novlan_6** with caution:

- Resetting VLAN settings can disrupt port communication (connection) with hosts or systems, including resetting the VLAN ID for an active iSCSI or IP partnership.
- You can reset a VLAN tag for a port that does not have VLAN tagging or does not have a configured IP address.

-storage yes | no

(Optional) Specifies whether an IPv4 address can be used for the backend storage attach function. The value *yes* indicates that this IPv4 address can be used for iSCSI target discovery and backend storage connectivity. You must specify *no* (default) if you are not using the storage attach IP address. If the IPv4 address associated with a specific port (on a node) is changed, the existing storage attach settings are retained. The values are *yes* and *no*.

-storage_6 yes | no

(Optional) Specifies whether an IPv6 address can be used for the backend storage attach function. The value *yes* (the default) indicates that this IPv6 address can be used for iSCSI target discovery and backend storage connectivity. You must specify *no* if you are not using a storage attach IP address. If the IPv6 address associated with a specific port (on a node) is changed, the existing storage attach settings are retained. The values are *yes* and *no*.

-force

(Optional) Forces an IP address change for a node Ethernet port even if this causes backend controllers being removed or MDisks going into a degraded or offline state.

Important: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information .

port_id

(Required) Specifies which port (1, 2, 3, or 4) to apply changes to.

Description

The **cfgportip** command either sets the IP address of an Ethernet port for iSCSI, or configures the MTU of a group of ports. 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 **chsystemip** command to assign clustered system IP addresses.

Remember: When IP addresses are configured with the same remote replication port group ID (for redundancy) to each node of an I/O group, make sure that the same Ethernet port for both nodes is used during configuration. MTU is set by using symmetric Ethernet ports from the same I/O group. To make sure alternative remote replication port groups work with the same MTU settings, symmetric Ethernet ports must be configured for remote replication port groups.

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.

If an IP address is specified for a host, the specified port can be discovered by hosts using the iSNS server (or other discovery mechanisms such as SendTargets). These IP addresses are not reported to partner systems in order to create TCP sessions used for remote copy. These ports can also not be used for login to and SendTargets based discovery of backend iSCSI Storage controllers.

IP addresses that are specified for remote copy cannot be discovered by hosts, which means they cannot be used for host attachment. These ports are not reported to partner systems in order to create TCP sessions for remote copy. These ports can also not be used to login to and for SendTargets when considering discovery of backend iSCSI Storage controllers.

IP addresses that are specified for storage cannot be discovered by hosts, which means they cannot be used for host attachment. These IP addresses are not reported to partner systems to create and set up TCP sessions for remote copy.

To use the same IP address for both host I/O and backend storage attach functions (but not for remote copy):

- The **-host** parameter must be set to *yes*.
- The **-storage** parameter must be set to *yes*.
- The **-remotecopy** parameter must be set to *no*.

In these instances, these IP addresses can be discovered by hosts. These IP addresses can also be used for backend storage controller discovery and login for iSCSI based migration and virtualization.

To use the same IP address for both backend storage attach functions and remote copy functions (but not for host I/O operations):

- The **-storage** parameter must be set to *yes*.
- The **-remotecopy** parameter must be specified with the required remote copy port group ID.
- The **-host** parameter must be set to *no*.

In such cases, these IP addresses can be used to discover and connect to backend iSCSI storage controllers. These IP addresses can also be used for IP-based remote copy.

To use the same IP address for both host I/O and remote copy functions (but not for backend storage attach functions):

- The **-host** parameter must be set to *yes*.
- The **-remotecopy** parameter must be invoked with the required remote copy port group ID.
- The **-storage** parameter must be set to *no*.

In such cases, these IP addresses can be discovered by hosts as well as used for IP-based Remote Copy but not for backend storage attach.

Use the **lsportip** command with the optional **ethernet_port_id** parameter to list the port IP addresses for the specified port.

Remember:

If **cfgportip** is used to modify the IP address that is associated with a specific Ethernet port without specifying a new VLAN ID, the new (modified) IP address inherits the existing VLAN ID setting of the earlier IP address (IPv4 or IPv6).

An invocation example for IPv4

```
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

```
cfgportip -node 1 -ip_6 3:3:0:4::0 -gw_6 ffe8::0 -prefix_6 64 2
```

The resulting output:

No feedback

An invocation example to set an MTU of 1600 on port 1 in I/O group 0

```
cfgportip -mtu 1600 -iogrp 0 1
```

The resulting output:

No feedback

An invocation example to set the MTU to its default value

```
cfgportip -defaultmtu -iogrp 0 1
```

The resulting output:

No feedback

An invocation example configuring a new IPv4 address for IP-based replication

```
cfgportip -node 1 -ip 9.8.7.1 -gw 9.0.0.1 -mask 255.255.255.0 -remotecopy 1 -host no -host_6 no 1
```

The resulting output:

No feedback

An invocation example configuring a new IPv4 address for host attach

```
cfgportip -node 1 -ip 9.8.7.1 -gw 9.0.0.1 -mask 255.255.255.0 -host yes 1
```

The resulting output:

No feedback

An invocation example configuring replication for an existing IPv6 address

```
cfgportip -node 1 -remotecopy_6 2 1
```

The resulting output:

No feedback

An invocation example configuring host attach for a new IPv6 address

```
cfgportip -node 1 -ip_6 2001:db8::1:0:0:1 -host_6 yes 1
```

The resulting output:

No feedback

An invocation example configuring a new IPv4 address with the VLAN ID 105

```
cfgportip -node 1 -ip 9.8.7.1 -gw 9.0.0.1 -mask 255.255.255.0 -vlan 105 1
```

The resulting output:

No feedback

An invocation example for configuring a new IPv6 address with the VLAN ID 1063

```
cfgportip -node 1 -ip_6 2001:db8::1:0:0:101 -prefix_6 64 -gw_6 2001:db8::1:0:0:1 -vlan_6 1063 1
```

The resulting output:

No feedback

An invocation example for configuring a new IPv4 address for the backend storage attach function using iSCSI

```
cfgportip -node 1 -ip 9.8.7.1 -gw 9.0.0.1 -mask 255.255.255.0 -storage yes -remotecopy 0 -host no 1  
cfgportip -node 1 -ip 9.8.7.1 -gw 9.0.0.1 -mask 255.255.255.0 -storage yes -host no 1
```

The resulting output:

No feedback

An invocation example for configuring a new IPv4 address for host attach only

```
cfgportip -node 1 -ip 9.8.7.1 -gw 9.0.0.1 -mask 255.255.255.0 -host yes -storage no 1  
cfgportip -node 1 -ip 9.8.7.1 -gw 9.0.0.1 -mask 255.255.255.0 -host yes -storage no -remotecopy 0 1  
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 configuring a new IPv4 address for IP-based replication

```
cfgportip -node 1 -ip 9.8.7.1 -gw 9.0.0.1 -mask 255.255.255.0 -storage no -remotecopy 1 -host no 1
```

The resulting output:

No feedback

An invocation example for configuring the storage attach function for a new IPv6 address

```
cfgportip -node 1 -ip_6 2001:db8::1:0:0:1 -storage_6 yes 1  
cfgportip -node 1 -ip_6 2001:db8::1:0:0:1 1
```

The resulting output:

No feedback

An invocation example for changing the storage specification for an existing IPv6 address

```
cfgportip -node 1 -storage_6 no 1  
cfgportip -node 1 -storage_6 yes 1
```

The resulting output:

No feedback

chbanner

Use the **chbanner** command to configure the login message that is displayed during CLI Secure Shell (SSH) login.

Syntax

```
➤ chbanner — [ -file file_path ] [ -enable ] [ -disable ] [ -clear ] ➤
```

Parameters

-file *file_path*

(Optional) Specifies the path to the file on the configuration node that contains the new login message.

-enable

(Optional) Enables the login message.

-disable

(Optional) Disables the login message.

-clear

(Optional) Clears the login message.

Description

This command configures the login message that is displayed during CLI SSH login. Use this for warnings or disclaimers or anything else you need to display in your login screen before logging in.

The file that contains the login message must be copied to the configuration node before you specify `chbanner -file`. If a configuration node failover occurs between copying the file that contains the login message and running the command, the temporary file must be copied to the new configuration node.

To set a login message that uses a SAN administrator's workstation:

1. Use a suitable text editor to create the message and save the file with a recognizable name.
2. Use a secure copy client to copy the file to the configuration node of the system to be configured.
3. Specify the management IP address of the system that is to be configured.
4. Log into the system to be configured.
5. Use the **chbanner** command to set the login message.

An invocation example

```
chbanner -file /tmp/loginmessage
```

The detailed resulting output:

No feedback

An invocation example

```
chbanner -enable
```

The detailed resulting output:

No feedback

An invocation example

```
chbanner -disable
```

The detailed resulting output:

No feedback

An invocation example

```
chbanner -clear
```

The detailed resulting output:

No feedback

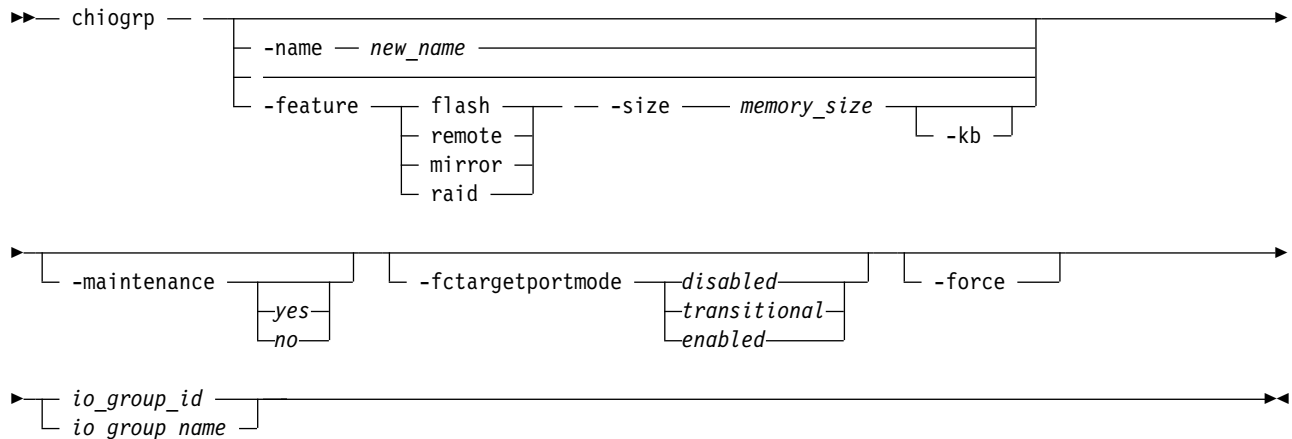
chcluster (Discontinued)

Attention: The **chcluster** command has been discontinued. Use the **chsystem** command instead.

chiogrp

Use the **chiogrp** command to modify the name of an I/O group, or the amount of memory that is available for RAID arrays, Copy Services, FlashCopy services, or volume 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 RAID arrays, Copy Services, or volume mirroring. You must specify this parameter with the **-size** parameter. You cannot specify this parameter with the **-name** parameter.

- *flash* specifies the amount of memory that is used for FlashCopy.
- *remote* specifies the amount of memory that is used for remote copy processing. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap.
- *mirror* specifies the amount of memory that is used for volume mirroring operations.
- *raid* specifies the amount of memory that is used for RAID arrays.

Note: Specifying **remote** changes the amount of memory that is available for remote copy processing. Any volume that is in a remote copy relationship uses memory in its I/O group, including master and auxiliary volumes, and volumes that are in inter-system or intra-system relationships.

-size *memory_size*

(Optional) Specifies the amount of memory that is available for the specified RAID arrays, Copy Services, or volume 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.

-maintenance *yes | no*

(Optional) Specifies whether the I/O group must be in maintenance mode. The I/O group must be placed in maintenance mode while performing service procedures on storage enclosures. After you enter maintenance mode, it continues until either:

- It is explicitly cleared
- Thirty minutes elapse

Note: Changing the maintenance mode on any I/O group changes the maintenance mode on all I/O groups.

-fctargetportmode *disabled | transitional | enabled*

(Optional) Specifies the Fibre Channel (FC) host port mode of the I/O group. The values are *disabled*, *transitional*, or *enabled*.

-force

(Optional) Specifies that an FC host port be disabled or enabled, even if disruption to host I/O might occur as a result. You can only specify **-force** with **-fctargetportmode**.

Important: Specifying **-force** might result in a loss of access. Use it only under the direction of the your product support information.

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.

Description

The **chiogrp** command modifies the name of an I/O group or the amount of memory that is available for RAID arrays, Copy Services, or volume mirroring.

Use the **-feature** and **-size** parameters (together) to change the amount of memory available in the I/O group to one of the following types:

- FlashCopy
- Volume mirroring
- RAID
- Remote copy, including Metro Mirror, Global Mirror, and HyperSwap.

For example:

```
chiogrp -feature flash -size 40 0
```

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 RAID arrays, Copy Services, or volume mirroring operations by specifying the **-feature** **flash | remote | mirror** parameter - and a memory size. For volume mirroring and Copy Services (Flash Copy®, Metro Mirror, Global Mirror, and HyperSwap), 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 amount of memory for FlashCopy is 20 MB.
- The default amount of memory for remote copy (which includes Metro Mirror, Global Mirror, and HyperSwap) is 20 MB.
- The default memory size for mirrored volumes is 20 MB.

- The default memory size for RAID arrays is 40 MB.
- The maximum amount of memory that can be specified for FlashCopy is 512 MB. For 64-bit systems, the maximum is 2048 MB.
- The maximum amount of memory for remote copy (which includes Metro Mirror, Global Mirror, and HyperSwap) is 512 MB.
- The maximum memory size that can be specified for mirrored volumes is 512 MB.
- The maximum memory size for RAID arrays is 512 MB.

The maximum combined amount of memory across all features is 552 MB.

Note: For 64-bit systems, the maximum is 2600 MB. Some systems running 64-bit mode might have 2 GB of bitmap space to use for FlashCopy, which is enough for 4 PB of data space to be used per I/O group. For example, Metro Mirror, Global Mirror, Volume Mirroring, and RAID share 552 MB of bitmap space, which is enough to use 1080 PB of data space per I/O group. Older systems, such as those running 32-bit code, might be subject to a 740 MB limit.

Table 26 demonstrates the amount of memory that is required for RAID arrays, Copy Services, and volume mirroring. Each 1 MB of memory provides the following volume capacities and grain sizes:

Table 26. Memory required for RAID arrays, Copy Services, and volume mirroring

Feature	Grain size	1 MB of memory provides the following volume capacity for the specified I/O group
Metro Mirror and Global Mirror	256 KB	2 TB of total Metro Mirror and Global Mirror volume capacity
HyperSwap	256 KB	2 TB of total HyperSwap volume capacity Note: For 2 TB of HyperSwap volume capacity, 1 MB must be assigned in each caching I/O group.
FlashCopy	256 KB	2 TB of total FlashCopy source volume capacity
FlashCopy	64 KB	512 GB of total FlashCopy source volume capacity
Incremental FlashCopy	256 KB	1 TB of total Incremental FlashCopy source volume capacity
Incremental FlashCopy	64 KB	256 GB of total Incremental FlashCopy source volume capacity
Volume mirroring	256 KB	2 TB of mirrored volumes

Table 27 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 27. RAID level comparisons

Level	Member count	Approximate capacity	Redundancy	Approximate bitmap memory cost
RAID-0	1-8	$MC * MS$	None	$(1 \text{ MB per } 2 \text{ TB of } MS) * MC$
RAID-1	2	MS	1	$(1 \text{ MB per } 2 \text{ TB of } MS) * (MC/2)$
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.

Table 27. RAID level comparisons (continued)

Level	Member count	Approximate capacity	Redundancy	Approximate bitmap memory cost
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 strip size for RAID-5 is ~1.15 MB for the first 2 TB of MS.

For multiple Flash Copy® 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 volume and a 16 GB target volume. Alternatively, for a mapping with a 256 KB grain size, 8 KB of memory allows two mappings between one 8 GB source volume and two 8 GB target volumes.

After you create a Flash Copy® mapping, if you specify an I/O group other than the I/O group of the source volume, the memory accounting goes towards the specified I/O group, not towards the I/O group of the source volume.

An invocation example to create a new I/O group testiogrpone

```
chiogrp -name testiogrpone io_grp0
```

The resulting output:

No feedback

An invocation example for changing the amount of Flash Copy® memory in io_grp0 to 30 MB

```
chiogrp -feature flash -size 30 io_grp0
```

The resulting output:

No feedback

An invocation example for changing the amount of RAID memory in I/O group 0 to 512 MB

```
chiogrp -feature raid -size 512 0
```

The resulting output:

No feedback

chnode


Use the **chnode** / **chnodecanister** command to change the name that is assigned to a node or node canister as well as other options. You can then use the new name when running subsequent commands. All parameters that are associated with this command are optional. However, you must specify one or more parameters.

Syntax

```

>>— chnode — | chnodecanister — [ -iscsialias — alias — ] [ -noiscsialias — ] [ -failover — ]

```



Parameters

-iscsialis *alias*

(Optional) Specifies the iSCSI name of the node or node canister. The maximum length is 79 characters.

Important: Do not use spaces for the iSCSI alias name.

-noiscsialias

(Optional) Clears any previously set iSCSI name for this node or node canister. 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 or node canister in the I/O group. When there is no partner node or node canister, the values set are applied to the partner node or node canister when it is added to the clustered system (system). If this parameter is used when there is a partner node or node canister, the name or alias of that node or node canister changes.

-name *new node or nodecanister name*

(Optional) Specifies the name to assign to the node or node canister.

Note: Node or node canister names supplied with `-name` on **chnode** / **chnodecanister** commands must not be in use already as node or node canister names or as node or node canister failover names.

Important: The iSCSI Qualified Name (IQN) for each node or node canister is generated using the clustered system and node or node canister names. If you are using the iSCSI protocol, changing either name also changes the IQN of all of the nodes or node canisters in the clustered system and might require reconfiguration of all iSCSI-attached hosts.

```
-identify yes | no
```

(Optional) Allows you to control the light-emitting diode (LED) used on the node. The values are *yes* or *no*.

```
-site site id | site name
```

(Optional) Specifies the numeric site value or site name for the existing node. The value is 1 or 2.

Note: The site assigned to the node cannot be changed if the system topology is HyperSwap or stretched.

-nosite

(Optional) Resets the site value.

```
object id | object name
```

(Required) Specifies the object name or ID that you want to modify. The variable that follows the parameter is either:

- The object name that you assigned when you added the node to the clustered system
- The object 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 or node canister. The name can then be used to identify the node or node canister in subsequent commands.

The **failover** parameter is used to specify values that are normally applied to the partner node or node canister in the I/O group. When the partner node or node canister is offline, the iSCSI alias and IQN are assigned to the remaining node or node canister in the I/O Group. The iSCSI host data access is then preserved. If the partner node or node canister is offline when these parameters are set, the node or node canister they are set on handles iSCSI I/O requests to the iSCSI alias specified, or the IQN that is created using the node or node canister name. If the partner node or node canister in the I/O group is online when these parameters are set, the partner node or node canister handles iSCSI requests to the iSCSI alias specified, and its node or node canister name and IQN change.

To change the name of the node (with I/O running) :

1. Make sure the host system has active sessions with both node canisters in the I/O group (hosting the volume on which the I/O occurs).
2. Change name of one node canister using **chnode** command.
3. From the host system, log out of the node canister whose name changes.
4. Rediscover the target iSCSI qualified name (IQN) from the host using the host operating system's discovery mechanism.
5. Login with the new target IQN discovered on the host system, and make sure the login succeeds.
6. Repeat steps 2-5 with the other node canister.

Note: When using VMware ESX, delete the static paths (in the iSCSI initiator properties) that contain the old target IQN.

This ensures that the node canister name change does not impact iSCSI I/O during events such as a target failover.

An invocation example

```
chnode -name newname -identify yes node8
```

The resulting output:

No feedback

An invocation example

```
chnode -name testnodeone nodeone
```

The resulting output:

No feedback

An invocation example

```
chnodecanister -name testnodeone nodeone
```

The resulting output:

No feedback

An invocation example

```
chnode -site 1 node2
```

The resulting output:

No feedback

An invocation example

```
chnodecanister -site 1 node2
```


The resulting output:

No feedback

chnodebattery

Use the **chnodebattery** command to set or clear the light-emitting diode (LED) on a hot-swappable battery (in a node). This command applies to SAN Volume Controller 2145-DH8 systems.

Syntax



Parameters

-identify

(Optional) Allows you to control the light-emitting diode (LED).

-battery *battery_id*

(Optional) Specifies the battery that is in the node.

-remove

(Optional) Specifies battery removal and terminates any calibration that runs on another battery.

node

(Optional) Specifies the node that the battery is in.

Description

This command notifies the battery back-up (BBU) driver that a user wants to remove a battery.

An invocation example to make the fault LED flash on battery 1 in node 3

```
chnodebattery -identify yes -battery 1 3
```

The resulting output:

No feedback

An invocation example to remove battery 1 in node 3

```
chnodebattery -remove -battery 1 3
```


The resulting output:

No feedback

chnodebootdrive

Use the **chnodebootdrive** command to change a drive or synchronize the drives on a system if there is a broken drive or field-replaceable unit (FRU) replacement drive. This command applies to SAN Volume Controller 2145-DH8 systems.

Syntax



Parameters

-sync

(Optional) Specifies synchronization of drives marked `can_sync`.

-force

(Optional) Forces synchronization (though taking the node offline might cause a volume to go offline).

Important: Using the `force` parameter might result in a loss of access. Use it only under the direction of your product support information.

-identify *yes* | *no*

(Optional) Allows you to control the operation of the light-emitting diode (LED) of the drive in the specified slot.

-slot *slot_id*

(Optional) Specifies the boot drive slot. It must be used with the **-identify** parameter.

resetleds

(Optional) Clears the identify LEDs of all drives in the specified node and indicates `-identify no` is specified.

node_id | *node_name*

(Optional) Specifies the ID or name of the node.

Description

The command identifies and synchronizes drive information for system drives.

Specifying **-sync** causes a node reboot on the specified node. This reboot is not successful if any volume is dependent on that node.

Important: If **-force** is also specified, the system does not check for dependent volumes.

An invocation example

```
chnodebootdrive
```

The resulting output:

```
No feedback
```

An invocation example

```
chnodebootdrive -identify yes -slot 1 1
```

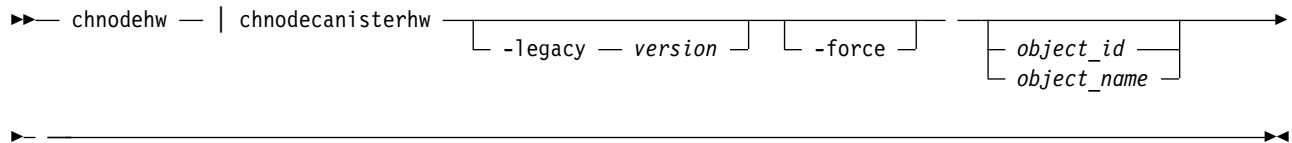
The resulting output:

```
No feedback
```

chnodehw (SVC) / chnodecanisterhw (Storwize family products)

Use the **chnodehw** / **chnodecanisterhw** command to update the hardware configuration for a node or node canister.

Syntax



Parameters

-legacy version

(Optional) Sets the hardware configuration to make it compatible with the 6.3.0.0 code level. The format is four decimal numbers separated by periods, and there can be up to sixteen characters.

-force

(Optional) Allow the node to restart and change its hardware configuration even if this will cause volumes to go offline.

Important: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information.

```
object id | object name
```

(Optional) Specifies the object name or ID.

Description

This command automatically reboots the node or node canister if the node or node canister hardware is different than its configured hardware. After rebooting, the node or node canister begins to use its hardware, and does not use the previous configuration.

Attention: When you run the **chnodehw** command to change the configured hardware for a node:

- Small Computer System Interface-3 (SCSI-3) reservations (through that node) are removed
- Small Computer System Interface-3 (SCSI-3) registrations (through that node) are removed

Use the **-legacy** parameter if you want to establish a partnership with another clustered system that is running an earlier level of code than the local system. The value supplied for the **-legacy** parameter must be the code level of the other clustered system.

An invocation example of how to update the node hardware configuration of node ID 7

chnodehw 7

The resulting output:

No feedback

An invocation example of how to update the node hardware configuration for the node named node7 (including if the node reboot causes an I/O outage)

```
chnodehw -force node7
```

The resulting output:

No feedback

An invocation example of how to update the node hardware configuration for compatibility with code level 6.3.0.0

```
chnodehw -legacy 6.3.0.0 node2
```

The resulting output:

No feedback

An invocation example of how to update the node canister hardware configuration of canister ID 7

```
chnodecanisterhw 7
```

The resulting output:

No feedback

An invocation example of how to update the node canister hardware configuration for canister7 (including if the canister reboot causes an I/O outage)

```
chnodecanisterhw -force canister7
```


The resulting output:

No feedback

chquorum

Use the **chquorum** command to change the quorum association.

Syntax



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* | *mdisk_name* | -drive *drive_id*

(Optional) Specifies the MDisk or drive to be this quorum ID.

Note: SAN Volume Controller uses MDisks only.

-override *yes|no*

Enables the automatic quorum selection to be overridden. In this state, the quorum disk is only moved if the resources are offline. Do not use this parameter unless a specific quorum disk is required for the configuration.

quorum_id

(Required) Specifies which quorum ID to change. Permitted values are values are 0, 1, and 2.

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 **lsquorum** command.

Remember: You cannot use this command to change the active quorum device when you use an IP quorum application. To change the active IP quorum application, the quorum application must be restarted. The quorum application that connects first is chosen and is active (if valid).

The **chquorum** command is not synchronous, but usually takes only a few seconds to complete. In some situations it can take several minutes.

The clustered system (system) uses the quorum disk or drive as a tie breaker when exactly half of the nodes that were previously a member of the system are present.

Attention: Only assign quorum disks 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 use of a quorum disk or drive allows the system to manage a SAN fault that splits the system exactly in half. One half of the system continues to operate and the other half stops until SAN connectivity is restored.

There is only one quorum disk or drive; however, the system uses three as quorum candidates. The system selects the actual quorum disk or drive from the pool of quorum candidates. The quorum candidates also hold a copy of important system 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 storage pool.

The number of extents required depends on the extent size for the storage pool containing the MDisk. Table 28 provides the number of extents reserved for quorum use by extent size.

Table 28. 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
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 system automatically assigns quorum indexes.

You can set the active quorum disk or drive with the **-active** parameter. This can be useful in a system configuration to ensure that the most highly-available quorum disk or drive is used.

Note: Quorum disks must be allocated one per site when the system topology is stretched or hyperswap.

An invocation example

```
chquorum -mdisk 45 2
```

The resulting output:

No feedback

chsecurity

Use the **chsecurity** command to change the Secure Sockets Layer (SSL) or Transport Layer Security (TLS) security settings for a clustered system (system).

Syntax

```
►►—chsecurity— —sslprotocol—security_level—————►►
```

Parameters

-sslprotocol *security_level*

(Required) Specifies the numeric value for the security level setting, which can take any value from 1 to 3.

A security level setting of:

- 1 disallows SSL 3.0
- 2 allows TLS 1.2 only
- 3 additionally disallows TLS 1.2 cipher suites that are not exclusive to 1.2

Description

This command changes the SSL or TLS security settings on a clustered system.

Important: If you use SSL or TLS, changing the security could disrupt these services.

If this occurs:

1. Wait 5 minutes and try again. (Wait for any services to restart.)
2. Confirm that the SSL or TLS implementation is up to date and supports the specified level of security.
3. If necessary, revert to an earlier version of SSL or TLS security.

An invocation example

```
chsecurity -sslprotocol 3
```

The resulting output:

Changing the security level could disable the GUI connection. If this happens use the CLI prompt to change to a known good level. Are you sure you wish to continue? (y/yes to confirm)

chsite

Use the **chsite** command to change the site name.

Syntax

```
►►—chsite— —name—new_site_name— —site_id—  
                                  existing_site_name—————►►
```

Parameters

-name *new_site_name*

(Required) Specifies the new or existing name for the site.

site_id | *existing_site_name*

(Required) Specifies the existing site ID or site name that is being changed.

Description

This command changes the site name.

Remember: This command is only applicable when a system is configured as a stretched system or a HyperSwap system (using the **chsystem -topology** command).

In a stretched configuration these applications are spread across two or more geographic locations or sites:

- Nodes
- Storage
- Host servers
- Infrastructure

An invocation example

```
chsite -name Quorum 3
```

The resulting output:

No feedback

chsystem

Use the **chsystem** command to modify the attributes of an existing clustered system (system). Enter this command any time after a system 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

```
➤➤ chsystem — — [ -name — system_name ] [ -rcbuffersize — new_size_in_MB ] —————➤
➤ [ -speed — fabric_speed ] [ -alias — id_alias ] —————➤
➤ [ -icatip — icat_console_ip_address ] [ -invemailinterval — interval ] —————➤
➤ [ -gmlinktolerance — link_tolerance ] [ -gmmxhostdelay — max_host_delay ] —————➤
➤ [ -icatip — ipv4_icat_ip_address ] [ -icatip_6 — ipv6_icat_ip_address ] —————➤
➤ [ -ntpip — ipv4_ntp_ip_address ] [ -ntpip_6 — ipv6_ntp_ip_address ] —————➤
➤ [ -isnsip — sns_server_address ] [ -isnsip_6 — ipv6_sns_server_address ] —————➤
```



Parameters

-name *system_name*

(Optional) Specifies a new name for the system.

Important: The Internet Small Computer System Interface (iSCSI) Qualified Name (IQN) for each node is generated by using the system and node names. If you are using the Internet Small Computer System Interface (iSCSI) protocol, changing either name also changes the IQN of all of the nodes in the system and might require reconfiguration of all iSCSI-attached hosts.

-rcbuffersize *new_size_in_MB*

(Optional) Specifies the amount of memory, in megabytes (MB), to use on each node for Metro Mirror and Global Mirror communications, from 48 to 512 MB. The default is 48 MB.

Important: Adjust this setting only when directed by your support team.

All nodes in the system must be online and have a minimum of 8 gigabytes (GB) - or 8192 megabytes (MB) - of memory to change this setting.

Remember:

- Before changing this setting you must stop all partnerships with this system.
- This parameter operates on the local system only and changing it is disruptive to mirroring operations.

-speed *fabric_speed*

(Optional) Specifies the speed of the fabric to which this system is attached. Valid values are 1Gb/s or 2Gb/s.

Attention: Changing the speed on a running system breaks input or output (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 alternative name that does not change the basic ID for the system, but does influence the VDisk_UID of every vdiskhostmap, both existing and new. These objects are created for a system whose ID matches the alias. Therefore, changing the system alias causes loss of host system access until each host scans for volumes that are presented by the system.

-invemailinterval *interval*

(Optional) Specifies the interval at which inventory emails are sent to the designated email recipients. The interval range is 0 to 15. The interval is measured in days. Setting the value to 0 turns off the inventory email notification function.

-gmlinktolerance *link_tolerance*

(Optional) Specifies the length of time, in seconds, for which an inadequate intersystem link is tolerated for a Global Mirror operation. The parameter accepts values from 20 to 86400 seconds in steps of 10 seconds. The default is 300 seconds. You can disable the link tolerance by entering a value of 0 for this parameter.

-gmmxhostdelay *max_host_delay*

(Optional) Specifies the maximum time delay, in milliseconds, at 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 volumes. You can use this parameter to increase the threshold from the default value of 5 milliseconds.

-icatip *icat_console_ip_address*

(Optional) Specifies the system's new IPv4 address that is used by the system. The format of this IP address must be a dotted decimal notation with the port; for example, 255.255.255.255:8080. If you specify this parameter, it overwrites any existing **-icatip_6** address.

-icatip_6 *icat_console_ipv6_address*

(Optional) Specifies the system's new IPv6 address. If you specify this parameter, it overwrites any existing **-icatip** address. The format of the IPv6 address must be:

- Eight colon-separated groups of four hexadecimal digits; for example:
[1234:1234:abcd:0123:0000:0000:7689:6576]:23
- Eight colon-separated groups of hexadecimal digits with leading zeros omitted; for example:
[1234:1234:abcd:123:0:0:7689:6576]:23
- Suppression of one or more consecutive all 0 groups; for example:
[1234:1234:abcd:123::7689:6576]:23

-ntpip *ipv4_ntp_ip_address*

(Optional) Specifies the IPv4 address for the Network Time Protocol (NTP) server. Configuring an NTP server address causes the system to use that NTP server as its time source. Specify the **-ntpip** parameter with a zero address to use another time source:

```
chsystem -ntpip 0.0.0.0
```

-ntpip_6 *ipv6_ntp_ip_address*

Note: Before you specify **-ntpip_6**, an IPv6 prefix and gateway must be set for the system. (Optional) Specifies the IPv6 address for the NTP server. Configuring an NTP server address causes the system to immediately start using that NTP server as its time source. To choose another time source, specify the **-ntpip_6** parameter with a zero address, as follows:

```
chsystem -ntpip_6 0::0
```

-isnsip *sns_server_address*

(Optional) Specifies the IPv4 address for the iSCSI storage name service (SNS). Specify the **-isnsip** parameter with a zero address to select another IPv4 iSCSI SNS server:

```
chsystem -isnsip 0.0.0.0
```

-isnsip_6 *ipv6_sns_server_address*

(Optional) Specifies the IPv6 address for the iSCSI SNS. Specify the **-isnsip_6** parameter with a zero address to select another configured IPv6 iSCSI SNS server:

```
chsystem -isnsip_6 0::0
```

-relationshipbandwidthlimit *bandwidth_in_mBps*

(Optional) Specifies the new background copy bandwidth in megabytes per second (MBps), from 1 to 1000. The default is 25 MBps. This parameter operates system-wide and defines the maximum background copy bandwidth that any relationship can adopt. The existing background copy bandwidth settings that are defined on a partnership continue to operate, with the lower of the partnership and volume rates attempted.

Note: Do not set this value higher than the default without establishing that the higher bandwidth can be sustained.

-infocenterurl *url*

Specifies the preferred online documentation URL to override the one used by the GUI. Because this information is interpreted by the Internet browser, the specified information might contain a hostname or an IP address.

Remember: View the currently configured URL in the GUI preferences window. This window can also help reset this value to the default setting.

-iscsiauthmethod *none* | *chap* -chapsecret *chap_secret*

(Optional) Sets the authentication method for the iSCSI communications of the system:

- *chap* indicates Internet Small Computer System Interface (iSCSI) authentication is turned on.

Remember: This turns on iSCSI partnership authentication whether a Challenge Handshake Authentication Protocol (CHAP) secret key is set for the system.

- *none* indicates that iSCSI partnership authentication is turned off.

-rcauthmethod *none* | *chap* -chapsecret *chap_secret*

(Optional) Turns authentication on or off for remote copy partnership requests that are native IP partnerships. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap. Additionally:

- *chap* indicates that remote copy authentication is turned on.

Remember: This action turns on authentication of remote copy partnership requests whether a Challenge Handshake Authentication Protocol (CHAP) secret key is set for the system.

- *none* indicates that remote copy partnership authentication is turned off.

-chapsecret *chap_secret*

(Optional) Sets the CHAP secret to be used to authenticate the system that uses iSCSI. This parameter is required if the **iscsiauthmethod *chap*** parameter is specified. The specified CHAP secret cannot begin or end with a space.

(Optional) Clears any previously set CHAP secret for iSCSI authentication. Then **nochapsecret** parameter cannot be specified whether **chapsecret** is specified.

(Optional) Specifies which layer a system is in. The system can create partnerships with systems in the same layer.

-cacheprefetch *on* | *off*

(Optional) Indicates whether cache prefetching is enabled or disabled across the system. Adjust this only when following direction from your product support information.

[illegible]

Remember: A partial mask (fewer than 64 characters) is zero-extended, meaning that any ports not specified are not enabled.

Specify the **lsportfc** command to display FC I/O port IDs.

[illegible]

Remember: A partial mask (fewer than 64 characters) is zero-extended, meaning that any ports not specified are not enabled.

Specify the **lsportfc** command to display FC I/O port IDs.

(Optional) Sets (or removes) High Temperature Mode (HTM). The values are on and off.

(Optional) Indicates the intended system topology, which is either *standard*, *stretched*, or *hyperswap*.

(Optional) Sets volume protection time (in minutes).

(Optional) Enables or disables volume protection. The values are yes and no.

(Optional) Enables or disables offloaded data transfers (ODX). The values are on and off.

(Optional) Enables Easy Tier and pool balancing acceleration. The values are on and off.

-maxreplicationdelay *value_in_seconds*

(Optional) Sets a maximum replication delay in seconds. The value must be a number from 0 to 360.

-partnershipexclusionthreshold *value_in_seconds*

(Optional) Sets the timeout for an I/O operation (in seconds) for remote systems. The value must be a number from 30 to 315 (default).

-ibmcustomer *customer_id*

(Optional) Specifies the customer number assigned when a software license that is automatically added to the entitlement database. The value must be a number that contains 7 - 10 digits (or blank).

-ibmcomponent *component_id*

(Optional) Specifies the component ID used for entitlement and call home system. The value is blank or SANVCNSW1.

-ibmcountry *country_id*

(Optional) Specifies the country ID used for entitlement and call home system. The value is blank or a 3-digit number.

Description

This command modifies specific features of a system. Multiple features can be changed by issuing a single command.

Using the **-ntpip** or **-ntpip_6** parameter allows the system to use an NTP server as an outside time source. The system 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 clock. In the NTP mode, the **setsystemtime** command is disabled.

All command parameters are optional, but you must specify at least one parameter.

Use the **chsystemip** command to modify the system IP address and service IP address.

Remember: Setting a CHAP secret key for the system does not turn on authentication for iSCSI hosts or remote copy partnerships. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap. Turn off authentication by issuing **-iscsiauthmethod** or **-rciauthmethod**.

The topology can only be set to hyperswap only if node sites have been configured. All nodes must have a site defined. If both nodes are defined in an I/O group they must be assigned to the same site (either 1 or 2; 3 cannot be used for nodes). You must have at least one I/O group with both nodes in site 1 and at least one I/O group with both nodes in site 2.

Note: If there are any active relationships defined, the system topology must be hyperswap.

An invocation example

```
chsystem -ntpip 9.20.165.16
```

The resulting output:

No feedback

An invocation example to set up an external NTP server

```
chsystem -ntpip 123.234.123.234
```

The resulting output:

No feedback

An invocation example to change the preferred infocenterurl value

```
chsystem -infocenterurl http://miscserver.company.com/ibm/infocenter
```

The resulting output:

No feedback

An invocation example to change the local port mask value

To set the local mask to sixty-two 0's and two 1's, indicating FC I/O ports with IDs 1 and 2 are capable of local node communication:

```
chsystem -localfcportmask 11
```

The resulting output:

No feedback

An invocation example to change the partner port mask value

To set the partner mask to sixty-three 0's and one 1, indicating that FC I/O port with ID 2 is capable of remote node communication:

```
svctask chsystem -partnerfcportmask 0010
```

The resulting output:

No feedback

An invocation example to change the HTM

```
chsystem -hightempmode on
```

The resulting output:

No feedback

An invocation example to set the system topology

```
chsystem -topology standard
```

The resulting output:

No feedback

An invocation example to set authentication for remote copy

```
chsystem -chapsecret ABCB1234 -iscsiauthmethod none -rcauthmethod chap
```

The resulting output:

No feedback

An invocation example to turn off volume protection

```
chsystem -vdiskprotectionenabled no
```

The resulting output:

No feedback

An invocation example to turn on volume protection and set the protection time to 60 minutes

```
chsystem -vdiskprotectionenabled yes -vdiskprotectiontime 60
```

The resulting output:

No feedback

An invocation example to turn on Easy Tier acceleration

```
chsystem -easytieracceleration on
```

The resulting output:

No feedback

An invocation example to turn on ODX

```
chsystem -odx on
```

The resulting output:

No feedback

An invocation example to set the maximum replication delay

```
chsystem -maxreplicationdelay 100
```

The resulting output:

No feedback

An invocation example to set the partnership exclusion threshold

```
chsystem -partnershipexclusionthreshold 120
```

The resulting output:

No feedback

An invocation example to specify an IBM customer ID, component ID, and country ID

```
chsystem -ibmcustomer 1928374 -ibmcomponent SANVCNSW1 -ibmcountry 001
```

The resulting output:

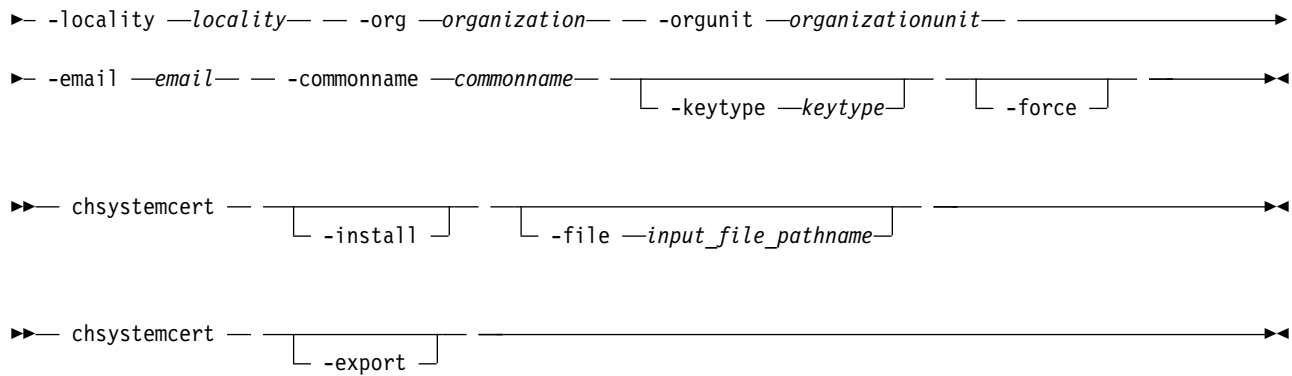
No feedback

chsystemcert

Use the **chsystemcert** command to manage the Secure Sockets Layer (SSL) certificate that is installed on a clustered system (system).

Syntax

```
➤ chsystemcert — [ -mkselfsigned ] [ -country country ] [ -state state ]
[ -locality locality ] [ -org organization ] [ -orgunit organizationunit ]
[ -email email ] [ -commonname commonname ] [ -keytype keytype ]
[ -validity days ]
➤ chsystemcert — [ -mkrequest ] [ -country country ] [ -state state ]
```



Parameters

-mkselfsigned

(Optional) Generates a self-signed SSL certificate. (If you do not specify **-mkselfsigned** you must specify **-mkrequest**, **-export**, or **-install**.)

-mkrequest

(Optional) Generates a certificate request. (If you do not specify **-mkrequest** you must specify **-mkselfsigned**, **-export**, or **-install**.)

-country *country*

(Optional for **-mkselfsigned** and required for **-mkrequest**) Specifies the 2-digit country code for the self-signed certificate or certificate request.

-state *state*

(Optional for **-mkselfsigned** and required for **-mkrequest**) Specifies the state information for the self-signed certificate or certificate request. The value can be an ASCII string from 0 - 128 characters.

-locality *locality*

(Optional for **-mkselfsigned** and required for **-mkrequest**) Specifies the locality information for the self-signed certificate or certificate request. The value can be an ASCII string from 0 - 128 characters.

-org *organization*

(Optional for **-mkselfsigned** and required for **-mkrequest**) Specifies the organization information for the SSL certificate. The value can be an ASCII string from 0 - 64 characters.

-orgunit *organizationunit*

(Optional for **-mkselfsigned** and required for **-mkrequest**) Specifies the organization unit information for the SSL certificate. The value can be an ASCII string from 0 - 64 characters.

-email *email*

(Optional for **-mkselfsigned** and required for **-mkrequest**) Specifies the email address that is used in the SSL certificate. The value can be an ASCII string from 0 - 64 characters.

-commonname *commonname*

(Optional for **-mkselfsigned** and required for **-mkrequest**) Specifies the common name for the SSL certificate. The value can be an ASCII string of 0 - 64 characters.

-validity *days*

(Optional) Specifies the number of days (1-9000) that the self-signed certificate is valid.

-keytype *keytype*

(Optional) Specifies the SSL certificate key type. The supported key types are:

- rsa2048
- ecdsa384
- ecdsa521

-install

(Optional) Installs a certificate. (If you do not specify **-install** you must specify **-mkselfsigned**, **-mkrequest**, or **-export**.)

-file

(Optional) Specifies the absolute path name of the certificate to install.

-export

(Optional) Exports the current SSL certificate. The certificate is exported to the `/dumps/certificate.pem` directory on the configuration node. (If you do not specify **-export** you must specify **-mkselfsigned**, **-mkrequest**, or **-install**.)

-force

(Optional) Specifies that the certificate request can be deleted.

Description

Use this command to manage the SSL certificate that is installed on a system. You can also

- Generate a new self-signed SSL certificate
- Create a certificate request to be copied from the system and signed by a certificate authority (CA)

Note: The signed certificate that is returned by the CA can be installed.

- Export the current SSL certificate (for example to allow the certificate to be imported into a key server)

Important: You must specify one of the following parameters:

- **-mkselfsigned**
- **-mkrequest**
- **-install**
- **-export**

An invocation example to create a self-signed certificate

```
chsystemcert -mkselfsigned
```

The detailed resulting output:

No feedback

An invocation example to create a self-signed certificate with a common name

```
chsystemcert -mkselfsigned -commonname weiland.snpp.com
```

The detailed resulting output:

No feedback

An invocation example to create a self-signed certificate with a key type and a 1-year validity period

```
chsystemcert -mkselfsigned -keytype ecDSA521 -validity 365
```

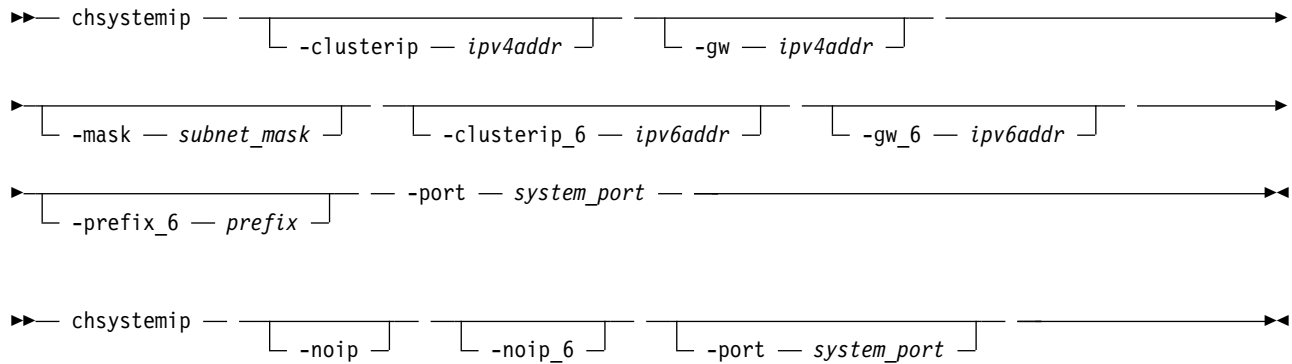
The detailed resulting output:

No feedback

chsystemip

Use the **chsystemip** command to modify the Internet Protocol (IP) configuration parameters for the clustered system (system).

Syntax



Parameters

-clusterip *ipv4addr*

(Optional) Changes the IPv4 system IP address. When you specify a new IP address for a system, the existing communication with the system is broken.

Important: The **-clusterip** parameter cannot be used if there are any active IPv4 partnerships with the system.

-gw *ipv4addr*

(Optional) Changes the IPv4 default gateway IP address of the system.

-mask *subnet_mask*

(Optional) Changes the IPv4 subnet mask of the system.

-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 system address for the port.

Important: The **-clusterip_6** parameter cannot be used if there are any active IPv6 partnerships with the system.

-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 *system_port*

(Required) 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 system. The first time you configure a second port, all IP information is required. Port 1 on the system must always have one stack fully configured.

There are two active system 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 system 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.

If there is no port 2 available on any of the system nodes, the **chsystemip** command fails.

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 system address. The configuration of port 2 is optional.

Service IP addresses for all ports and stacks are initialized to Dynamic Host Configuration Protocol (DHCP). A service IP address is always configured.

Note: If the **console_ip** is the same as IP address system port 1, Internet Protocol Version 4 (IPv4) followed by IPv6, change the **console_ip** when the system IP is changed. If the **console_ip** differs from the system port 1 IP address, do not change the **console_ip** when the system IP is changed.

To modify an IP address, list the IP address of the system by issuing the **lssystem** command. Modify the IP address by issuing the **chsystemip** command. You can either specify a static IP address or have the system assign a dynamic IP address.

Table 29 provides IP address formats that are supported.

Table 29. IP address list formats

IP type	IP address list format
IPv4	1.2.3.4
Full IPv6	1234:1234:abcd:0123:0000:0000:7689:6576
Full IPv6, leading zeros suppressed	1234:1234:abcd:123:0:0:7689:6576
IPv6 with zero compression	1234:1234:abcd:123::7689:6576

An invocation example

```
chsystemip -clusterip 9.20.136.5 -gw 9.20.136.1 -mask 255.255.255.0 -port 1
```

The resulting output:

No feedback

An invocation example

```
chsystemip -clusterip_6 2001:0db8:85a3:0000:0000:8a2e:0370:7334 -gw_6 2001:0db8:85a3:0000:0000:8a2e:0370:7334 -prefix_6 64 -port 2
```


The resulting output:

No feedback

chthrottle

Use the **chthrottle** command to change attributes associated with a specified throttle object.

Syntax



Parameters

-bandwidth *bandwidth_limit_in_mb*

(Optional) Specifies the bandwidth in MBps. This must be a numeric value 0 - 268435456.

Note: No bandwidth limit is set unless you specify this keyword.

-iops *iops_limit*

(Optional) Specifies the I/O operations limit. This must be a numeric value 0 - 33554432.

Note: No I/O operations limit is set unless you specify this keyword.

-name *throttle_name*

(Optional) Specifies the throttling object's name. This value must be an alphanumeric string up to 63 characters long.

throttle_id | *throttle_name*

(Required) Specifies the volume ID or name of the volume to throttle. The value must be a numeric or alphanumeric string up to 15 characters long.

Description

This command changes attributes associated with a specified throttle object.

An invocation example for changing the bandwidth limit to 100 for an offloaded throttle

```
chthrottle -bandwidth 100 offloadThrottle
```

The detailed resulting output:

No feedback

An invocation example with no throttling bandwidth limit specified for ID 0

```
chthrottle -bandwidth 100 0
```


The detailed resulting output:

No feedback

cleardumps

Use the **cleardumps** command to clear (or delete) the various dump directories on a specified node or node canister.

Syntax

►► `cleardumps` — `-prefix` — *directory_or_file_filter* — 

Parameters

-prefix *directory_or_file_filter*

(Required) Specifies the directory, files, or both to be cleared. If a directory is specified, with no file filter, all relevant dump or log files in that directory are cleared. You can use the following directory arguments (filters):

- `/dumps` (clears all files in all subdirectories)
- `/dumps/cimom`
- `/dumps/cloud`
- `/dumps/configs`
- `/dumps/easytier`
- `/dumps/elog`
- `/dumps/feature`
- `/dumps/iostats`
- `/dumps/iotrace`
- `/dumps/mdisk`
- `/home/admin/update`

In addition to the directory, you can specify a filter file. For example, if you specify `/dumps/elog/*.txt`, all files in the `/dumps/elog` directory that end in `.txt` are cleared.

Note: The following rules apply to the use of wildcards when using the 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:

```
>cleardumps -prefix "/dumps/elog/*.txt"
```

node_id | *node_name*

(Optional) Specifies the node to be cleared. 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 clustered system (system)
- 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 cleared.

You can clear all the dumps directories by specifying `/dumps` as the directory variable.

You can clear 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 `lsxxxxdumps` commands.

You can use this command to clear 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 cleared:

- `*svc.config*`
- `*.trc`
- `*.trc.old`

An invocation example

```
cleardumps -prefix /dumps/configs
```

The resulting output:

No feedback

An invocation example

```
cleardumps -prefix /dumps/easytier node_2
```

The resulting output:

No feedback

cpdumps

Use the **cpdumps** command to copy 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

```
➤— cpdumps — — -prefix ———┐ directory ———┐—————➤
                             └─ file_filter —┘   └─ node_name
                                                    └─ 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/cloud`
- `/dumps/configs`
- `/dumps/elogs`
- `/dumps/easytier`
- `(Storwize V7000)/dumps/enclosure`
- `/dumps/feature`
- `/dumps/iostats`
- `/dumps/iotrace`

- /dumps/mdisk
- /home/admin/update

In addition to the directory, you can specify a file filter. For example, if you specified /dumps/elog/*.txt, all files in the /dumps/elog 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:

```
>cleardumps -prefix "/dumps/elog/*.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 clustered system (system)
- 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, system 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 **lsdumps** command. You can track the status of a copy using the **lscopystatus** command.

An invocation example

```
cpdumps -prefix /dumps/configs nodeone
```

The resulting output:

No feedback

An invocation example

```
cpdumps -prefix /dumps/easytier node_2
```

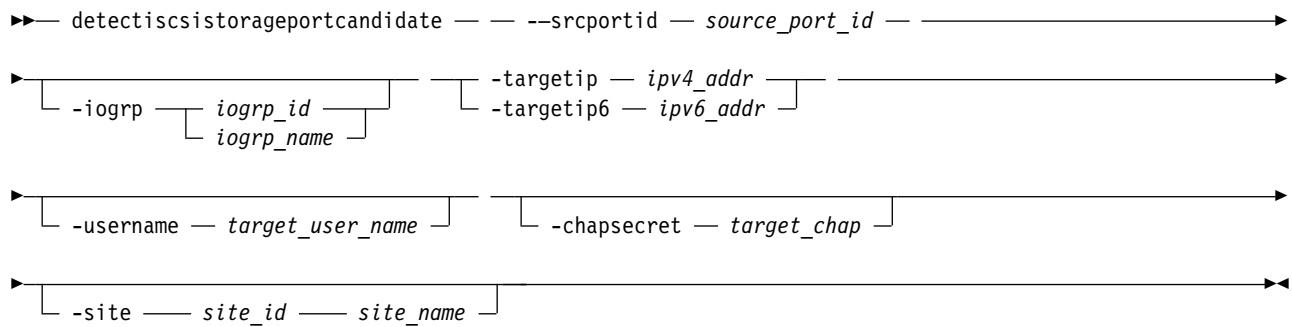
The resulting output:

No feedback

detectiscsistorageportcandidate

Use the **detectiscsistorageportcandidate** command to establish Internet Small Computer Systems Interface (iSCSI) login sessions from any nodes in a specified I/O group to a discovered backend iSCSI target controller.

Syntax



Parameters

-srcportid *source_port_id*

(Required) Specifies the source Ethernet port ID (indicated in the **lsportip** output) used to complete target controller discovery. The value must be a number 1 - 8.

- If you also specify **-iogrp**, you trigger discovery through the Ethernet port by using the *source_port_id* specified on all nodes in the I/O group.
- If you do not specify **-iogrp**, you trigger discovery through the Ethernet port by using the *source_port_id* specified on all nodes in the clustered system (system).

-iogrp *iogrp_id* | *iogrp_name*

(Optional) Specifies I/O group ID or name being detected. The *iogrp_id* value must be 0, 1, 2, or 3. The *iogrp_name* value must be an alphanumeric string.

If you specify this parameter, you trigger discovery through the *source_port_id* of both nodes for the specified I/O group. If no value is specified, discovery is triggered by using the *source_port_id* of all nodes in the system.

-targetip *ipv4_addr*

(Required if you do not specify **-targetip6**) Specifies the target iSCSI controller IPv4 address being detected that receives target discovery requests by using the *source_port_id* for all nodes in the specified I/O group.

-targetip6 *ipv6_addr*

(Required if you do not specify **-targetip**) Specifies the target iSCSI controller IPv6 address being detected that receives target discovery requests by using the *source_port_id* for all nodes in the specified I/O group.

-username *target_user_name*

(Optional) Specifies the target controller user name being detected. The value must be an alphanumeric string up to 256 characters.

If the target controller requires a *target_user_name* and *target_chap* for discovery, use the target's user name and chap to discover target controller.

Some controllers might require that you use the iSCSI qualified name (IQN) user name for discovery. Each nodes IQN is picked up automatically and used if required.

-chapsecret *target_chap*

(Optional) Specifies the *target_chap* required for discovery of the target iSCSI controller that is being detected. The value must be an alphanumeric string (case-sensitive) up to 79 characters.

-site *site_id* | *site_name*

(Optional) Specifies the site ID or site name of the target iSCSI controller that is being detected or discovered. If no I/O group is specified, discovery is done from all nodes present in the specified site

and if an I/O group is specified discovery is done only from a node that is part of the specified site and I/O group. The site ID must be 1 (the default) or 2. The site name must be an alphanumeric value.

Important: This parameter must be specified for a HyperSwap or stretched system.

The stretched system topology distributes the I/O group information into each site. Each node in the I/O group is associated with a different site. The back end storage controller of one site is visible only to nodes within that same site. As a result, you attempt storage controller discovery only from nodes that are in the same site.

Description

This command detects iSCSI backend storage controllers for migration and virtualization. This command helps with target iSCSI controller discovery. Use the **lsiscsistorageportcandidate** command to list information about the discovered candidate iSCSI target controller(s).

The target data is available until either another discovery is completed or the system undergoes a recovery procedure, which clears the previous discovery data. The command completes when either discovery from all source nodes completes or the command fails.

Use the **addiscsistorageport** command to establish sessions from any nodes in a specified I/O group to a discovered backend iSCSI controller. Use the **lsiscsistorageportcandidate** command to list information about the candidate iSCSI target controller.

Showing target discovery using an IPv4 IP address with a target iSCSI controller and source port ID 1

```
detectiscsstorageportcandidate -srcportid 1 -targetip IP1
```

The detailed resulting output:

No feedback

Showing target discovery from I/O group 3 using an IPv6 address with a target iSCSI controller and source port ID 1

```
detectiscsistorageportcandidate -srcportid 3 -targetip6 IP3 -iogrp 3 -username delluser -chapsecret password1
```

The detailed resulting output:

No feedback


dumpconfig (Discontinued)

Attention: The **dumpconfig** command is discontinued.

help

Use the **help** (or **man**) command to display help information for system commands.

Syntax



Parameters

command_name

(Optional) Indicates the command name.

Description

Use this command to display help information for system commands. If you specify a command name using *command_name*, the complete help file text for the command is displayed. If you do not specify a command name, a comprehensive list of all commands is displayed (with one brief descriptive line). This list includes these commands:

- `satask`
- `sainfo`
- `svcconfig`
- `svc_snap`
- `svc_livedump`

Remember: The `help` command alias is `man`.

An invocation example

`help`

The resulting output:

```
addhostiogr - Maps I/O groups to a host object.
addhostport - Adds worldwide port names (WWPNs) or iSCSI names to a host object.
addmdisk - Adds managed disks to a storage pool.
addnode - Adds a new (candidate) node canister to an existing system.
...
```

Isclustercandidate (Discontinued)

Attention: The `isclustercandidate` command has been discontinued. Use the `ispartnershipcandidate` command instead.

Iscluster (Discontinued)

Attention: The `iscluster` command is discontinued. Use a combination of the `ispartnership`, `ispartnershipcandidate`, and `isssystem` commands instead.

Isclusterip (Discontinued)

Attention: The `isclusterip` command has been discontinued. Use the `isssystemip` command instead.

Isclusterstats (Discontinued)


Attention: The `isclusterstats` command is discontinued. Use the `isssystemstats` command instead.

Isdiscoverystatus

Use the `Isdiscoverystatus` command to determine whether a discovery operation is in progress.

Syntax

```
►► isdiscoverystatus — [ -filtervalue — attribute_value ] [ -filtervalue? ] —►
```



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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):


```
lsdiscoverystatus -filtervalue "IO_group_name=io*"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- IO_group_id
- IO_group_name
- scope
- status

-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 **lsdiscoverystatus** command. The *fc_fabric* line will only appear if there is at least one Fibre Channel controller.

An invocation example

```
lsdiscoverystatus -delim :
```

The resulting output:

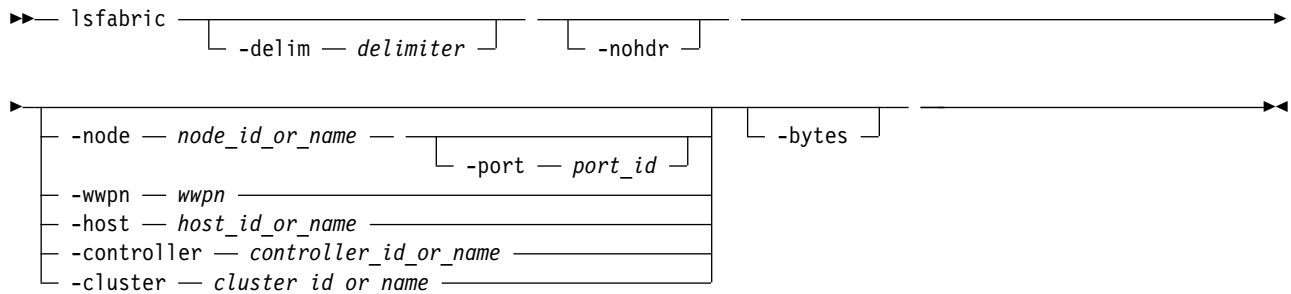
```
id:scope:I0_group_id:I0_group_name:status
0:fc_fabric:::active
1:sas_iogrp:0:io_grp0:inactive
3:sas_iogrp:2:io_grp2:active

6:iscsi:::inactive
```

lsfabric

Use the **lsfabric** command to generate a report that displays the Fibre Channel (FC) connectivity between nodes, controllers, and hosts.

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.

-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 worldwide port names (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 from a minimum of one through a maximum equal to the number of node Fibre Channel (FC) I/O ports. It 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 **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 clustered system (system) 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 **lsfabric wwpn wwpn** command for every configured WWPN of the specified controller. For example, a controller with 4 ports connected to a 8 node system with 2 counterpart SANs produces 64 lines of output.

-cluster *cluster_id_or_name*

(Optional) Specifies a system 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 **lsfabric wwpn wwpn** command for every known WWPN in the specified system. Output is sorted by remote WWPNs and then system WWPNs. This parameter can be used to check the state of connections within the local system or between the local and remote system. When the local system ID or name is specified, each node-to-node connection is listed twice: once from each end. For example, an eight-node system 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 system must be configured in a remote copy partnership with the local system. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap. It must appear in the **lssystem** view.

-bytes

(Optional) Displays all capacities as bytes.

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.

Remember: The value of the `local_port` field is the number of the node's Fibre Channel (FC) port.

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

Small Computer System Interface (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.

state blocked

This shows connections that are blocked due to the system's port mask settings.

type One of the following values is displayed:

- host
- node
- controller
- unknown
- nas

You can issue this command to view all the information about the connections that are available to your system.

Remember: The **lsfabric** command is limited to displaying 16,384 entries. If you have a large system configuration that exceeds these limits you must filter the output (for example, by node or node port) to view all fabric login records.

An invocation example using a delimiter

```
lsfabric -delim :
```

The resulting output, in which 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
```

An invocation example that shows unused (because the system's mask settings are blocked) node logins

```
lsfabric -delim :
```

The resulting output:

```
remote_wwpn:remote_nportid:id:node_name:local_wwpn:local_port:local_nportid:state:name:cluster_name:type
500507680304D190:021700:5:nodeA:500507680304A100:1:020300:active:node4:Cluster_9.115.2:node
500507680304D190:021700:2:nodeB:500507680308A101:2:021800:active:node4:Cluster_9.115.2:node
500507680304D190:021700:3:nodeC:500507680308190D:2:020A00:active:node4:Cluster_9.115.2:node
500507680308D190:011700:5:nodeA:500507680308A100:2:011000:blocked:node4:Cluster_9.115.2:node
500507680308D190:011700:2:nodeB:500507680304A101:1:010D00:blocked:node4:Cluster_9.115.2:node
500507680308D190:011700:3:nodeC:500507680304190D:1:011200:blocked:node4:Cluster_9.115.2:node
```

An invocation example that shows internal connectivity for node 1

```
lsfabric -internal -delim : -node 1
```

The resulting output:

```
remote_wwpn:remote_nportid:id:node_name:local_wwpn:local_port:local_nportid:state:name:cluster_name:type
500507680C520034:010E00:1:node1:500507680C210033:5:010200:active:node2:Cluster_9.19.88:node
500507680C520034:010E00:1:node1:500507680C220033:6:010F00:active:node2:Cluster_9.19.88:node
500507680C520034:010E00:1:node1:500507680C510033:9:010A00:active:node2:Cluster_9.19.88:node
500507680C520034:010E00:1:node1:500507680C520033:10:010B00:active:node2:Cluster_9.19.88:node
500507605EBFEA91:010900:1:node1:500507680C210033:5:010200:active:::expansion
500507605EBFEA91:010900:1:node1:500507680C220033:6:010F00:active:::expansion
500507605EBFEA91:010900:1:node1:500507680C510033:9:010A00:active:::expansion
500507605EBFEA91:010900:1:node1:500507680C520033:10:010B00:active:::expansion
500507605E828601:010100:1:node1:500507680C210033:5:010200:active:::expansion
500507605E828601:010100:1:node1:500507680C220033:6:010F00:active:::expansion
500507605E828601:010100:1:node1:500507680C510033:9:010A00:active:::expansion
500507605E828601:010100:1:node1:500507680C520033:10:010B00:active:::expansion
500507605E828611:010700:1:node1:500507680C210033:5:010200:active:::expansion
500507605E828611:010700:1:node1:500507680C220033:6:010F00:active:::expansion
500507605E828611:010700:1:node1:500507680C510033:9:010A00:active:::expansion
500507605E828611:010700:1:node1:500507680C520033:10:010B00:active:::expansion
500507680C210034:010000:1:node1:500507680C210033:5:010200:active:node2:Cluster_9.19.88:node
500507680C210034:010000:1:node1:500507680C220033:6:010F00:active:node2:Cluster_9.19.88:node
```

```

500507680C210034:010000:1:node1:500507680C510033:9:010A00:active:node2:Cluster_9.19.88:node
500507680C210034:010000:1:node1:500507680C520033:10:010B00:active:node2:Cluster_9.19.88:node
500507605EBFEAB1:010400:1:node1:500507680C210033:5:010200:active:::expansion
500507605EBFEAB1:010400:1:node1:500507680C220033:6:010F00:active:::expansion
500507605EBFEAB1:010400:1:node1:500507680C510033:9:010A00:active:::expansion
500507605EBFEAB1:010400:1:node1:500507680C520033:10:010B00:active:::expansion
500507680C510034:010D00:1:node1:500507680C210033:5:010200:active:node2:Cluster_9.19.88:node
500507680C510034:010D00:1:node1:500507680C220033:6:010F00:active:node2:Cluster_9.19.88:node
500507680C510034:010D00:1:node1:500507680C510033:9:010A00:active:node2:Cluster_9.19.88:node
500507680C510034:010D00:1:node1:500507680C520033:10:010B00:active:node2:Cluster_9.19.88:node
500507605EBFEA82:010500:1:node1:500507680C210033:5:010200:active:::expansion
500507605EBFEA82:010500:1:node1:500507680C220033:6:010F00:active:::expansion
500507605EBFEA82:010500:1:node1:500507680C510033:9:010A00:active:::expansion
500507605EBFEA82:010500:1:node1:500507680C520033:10:010B00:active:::expansion
500507605EBFEAA2:010600:1:node1:500507680C210033:5:010200:active:::expansion
500507605EBFEAA2:010600:1:node1:500507680C220033:6:010F00:active:::expansion
500507605EBFEAA2:010600:1:node1:500507680C510033:9:010A00:active:::expansion
500507605EBFEAA2:010600:1:node1:500507680C520033:10:010B00:active:::expansion
500507680C220034:010C00:1:node1:500507680C210033:5:010200:active:node2:Cluster_9.19.88:node
500507680C220034:010C00:1:node1:500507680C220033:6:010F00:active:node2:Cluster_9.19.88:node
500507680C220034:010C00:1:node1:500507680C510033:9:010A00:active:node2:Cluster_9.19.88:node
500507680C220034:010C00:1:node1:500507680C520033:10:010B00:active:node2:Cluster_9.19.88:node
500507605E828631:010800:1:node1:500507680C210033:5:010200:active:::expansion
500507605E828631:010800:1:node1:500507680C220033:6:010F00:active:::expansion
500507605E828631:010800:1:node1:500507680C510033:9:010A00:active:::expansion
500507605E828631:010800:1:node1:500507680C520033:10:010B00:active:::expansion
500507605E828621:010300:1:node1:500507680C210033:5:010200:active:::expansion
500507605E828621:010300:1:node1:500507680C220033:6:010F00:active:::expansion
500507605E828621:010300:1:node1:500507680C510033:9:010A00:active:::expansion
500507605E828621:010300:1:node1:500507680C520033:10:010B00:active:::expansion

```

Isfcportcandidate

Use the **Isfcportcandidate** command to list the Fibre Channel (FC) ports. This information is used to find open FC ports.

Syntax

```

>> Isfcportcandidate — [ -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 FC ports.

Note: The **lsfcportcandidate** command presents a list of host FC ports that are logged in to nodes. However, there are situations when the information that is presented might include host FC ports that are no longer logged in or even part of the SAN fabric. For example, if a host FC port is unplugged from a switch but **lsfcportcandidate** shows the worldwide port name (WWPN) that is logged in to all nodes, the incorrect entry is removed when another device is plugged in to the same switch port that previously contained the removed host FC port.

Table 30 shows the possible output:

Table 30. *lsfcportcandidate* output

Attribute	Description
fc_WWPN	Indicates the FC WWPN is logged in but unconfigured (not assigned to a host). This value must be 16 hexadecimal characters.

An invocation example

```
lsfcportcandidate
```


The resulting output:

```
fc_WWPN
200600A0B813B7AC
200600A0B813B7AD
```

lscsistorageport

Use the **lscsistorageport** command to display the details of the iSCSI login session(s) that are established from the Initiator's Internet Small Computer Systems Interface (iSCSI) source ports to iSCSI backend target controller ports.

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.

-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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.

- When you use a wildcard, enclose the filter entry within double quotation marks (""): `lsiscsistorageport -filtervalue id="1*"`

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- id
- status

-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 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 1-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.

row_id

(Optional) Specifies the row ID view and denotes the sessions that are established from the specified initiator node ports to a backend controller target iSCSI qualified name (IQN) through a target Internet Protocol (IP) address. The value must be a number 0 - 1024.

Description

This command displays the details of sessions that are established from the Initiator's iSCSI source ports to iSCSI backend target controller ports.

This table provides the attribute values that can be displayed as output view data.

Table 31. lsiscsistorageport output

Attribute	Description
id	Indicates the object identifier for any sessions from any clustered system (system) nodes to the iSCSI backend controller iSCSI qualified name (IQN) through an iSCSI backend controller target IP. The value must be a number 0 - 1023.
src_port_id	Indicates the source port identifier for the node Ethernet port number that is displayed in the lsportip output. The value is a number 0 - 7.
target_ipv4	Indicates the IPv4 address of the iSCSI backend controller target port that establishes a session from the initiator source port that is identified by the source port ID. The default value is blank.
target_ipv6	Indicates the IPv6 address of the iSCSI backend controller target port that establishes a session from the initiator source port that is identified by the source port ID. The default value is blank.
target_iscsiname	Indicates the IQN of the iSCSI backend controller target that establishes a session. The value must be an alphanumeric string of no more than 256 characters. The default value is blank.
controller_id	Indicates the controller ID that is displayed in the lscontroller output. The value must be a number 0 - 1023. The default value is 1024.

Table 31. *Isiscsistorageport* output (continued)

Attribute	Description
iogroup_list	<p>Indicates a colon-separated list of discovery result codes: The value must be an alphanumeric string of up to 32 characters. This field cannot be blank. The values are 0 and 1:</p> <ul style="list-style-type: none"> 0 indicates that the I/O group is available in the system, but discovery is either not triggered through the I/O group or discovery through the I/O group fails 1 indicates that the I/O group is present and discovery is successful through the I/O group <p>Note: The value - (dash) indicates that the I/O group is not valid or is not present in the system.</p>
status	<p>Indicates the connectivity status from all nodes in the system to the target port. The values are:</p> <ul style="list-style-type: none"> full If you specify a single I/O group by using the addiscsistorageport command and you establish the session from all nodes in the specified I/O group, the status is full. partial If you specify a single I/O group by using the addiscsistorageport command and you establish the session from a single node in the specified I/O group, the status is partial. none If you specify a single I/O group by using the addiscsistorageport command and you do not establish the session from any node in the specified I/O group, the status is none. <p>There is no default value. This field cannot be blank.</p>
connected	Indicates whether the established connection is from a specified Ethernet port of a target IQN and IP address. The values are yes and no.
site_id	Indicates the site ID (if the nodes being discovered belong to a site). This applies to stretched and HyperSwap systems.
site_name	Indicates the site name (if the nodes being discovered belong to a site). This applies to stretched and HyperSwap systems.
node_id	Indicates the node ID of the initiator node that establishes the session. The value must be a numeric value.
node_name	Indicates the node name of the initiator node that establishes the session. The value must be an alphanumeric string of no more than 16 characters (the default value is blank).
src_ipv4	Indicates the IPv4 IP address of the source port ID on a specified node. The default value is blank.
src_ipv6	Indicates the IPv6 IP address of the source port ID on a specified node. The default value is blank.
src_iscsiname	Indicates the IQN of the source node for which connectivity is being displayed for the target port. The value must be an alphanumeric string of no more than 256 characters (the default value is blank).

Before you specifying **lsiscsistorageport** such as in the examples below, you must:

- complete target discovery by using an IPv4 IP address of a target ISCSI controller through source port ID 0:
`detectiscsistorageportcandidate -targetip IP1 -srcportid 2`
- You would then specify session establishment by using **addiscsistorageport** for discovery output row 0 through I/O group 1:

```
addiscsstorageport -iogrp 1 0
```

3. Specify **lsiscsistorageport** to view the output (no *tgt_user_name* or *target_chap* is required for discovery or session establishment).

Specify **rmiscsistorageport** to remove a session.

A concise invocation example

```
lsiscsistorageport
```

The resulting output:

id	src_port_id	target_ipv4	target_ipv6	target_iscsiname	controller_id	controller_name	iogroup_li
0	4	192.168.82.90		iqn.1986-03.com.ibm:2145.temp.node1	3	controller3	0:1:-:-

A detailed invocation example

```
lsiscsistorageport 0
```

The resulting output:

```
id 0
src_port_id 4
target_ipv4 192.168.82.90
target_ipv6
target_iscsiname iqn.1986-03.com.ibm:2145.temp.node1
controller_id 0
iogroup_list 1:1:-:-
status full
site_id
site_name
node_id 17
node_name node1
src_ipv4 192.168.82.80
src_ipv6
src_iscsiname iqn.1986-03.com.ibm:2145.iscsicluster.node1
connected yes
node_id 20
node_name node2
src_ipv4 192.168.82.81
src_ipv6
src_iscsiname iqn.1986-03.com.ibm:2145.iscsicluster.node2
connected yes
node_id 16
node_name node3
src_ipv4 192.168.82.82
src_ipv6
src_iscsiname iqn.1986-03.com.ibm:2145.iscsicluster.node3
connected yes
node_id 18
node_name node4
src_ipv4 192.168.82.83
src_ipv6
src_iscsiname iqn.1986-03.com.ibm:2145.iscsicluster.node4
connected yes
```

lsiscsistorageportcandidate

Use the **lsiscsistorageportcandidate** command to list information about the candidate Internet Small Computer Systems Interface (iSCSI) target controller iSCSI Qualified Name (IQN) that is specified with the target IP from the specified initiator source ports.

Syntax

►► `lsiscsistorageportcandidate` — `-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 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 1-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 lists information about the last invocation of the **`detectiscsistorageportcandidate`** command. This command also lists the candidate iSCSI target controller IQNs that are visible at the specified target IP from the specified initiator ports along with indication of which initiator ports can see each discovered iSCSI target IQN.

This table provides the attribute values that can be displayed as output view data.

Table 32. *lsiscsistorageportcandidate* output

Attribute	Description
id	Indicates the row ID for the discovery output. There is one row per iSCSI qualified name (IQN) that can be viewed using the target IP address specified in <code>detectiscsistorageportcandidate</code> .
src_port_id	Indicates the source port ID. The value is a number 1 - 8.
target_ip	Indicates the target IPv4 address.
target_ip6	Indicates the target IPv6 address.
target_iscsiname	Indicates the (discovered) IQN that uses the target controller. The value is an alphanumeric string that is 256 characters long.
iogroup_list	<p>Indicates a colon-separated list of discovery result codes: The value must be an alphanumeric string of up to 32 characters. This field cannot be blank. The values are 0 and 1:</p> <ul style="list-style-type: none">• 0 indicates that the I/O group is available in the system, but discovery is either not triggered through the I/O group or discovery through the I/O group fails• 1 indicates that the I/O group is present and discovery is successful through the I/O group <p>Note: The value - (dash) indicates that the I/O group is not valid or is not present in the system.</p>

Table 32. *Isiscsistorageportcandidate* output (continued)

Attribute	Description
status	Indicates whether discovery was successful. The values are: <ul style="list-style-type: none"> • full • partial • none
configured	Indicates whether the discovered target IQN has any established sessions with source ports or target ports. The values are yes and no (default).
site_id	Indicates the site ID (if the nodes being discovered belong to a site). This applies to stretched and HyperSwap systems.
site_name	Indicates the site name (if the nodes being discovered belong to a site). This applies to stretched and HyperSwap systems.

An invocation example

First you must specify target discovery using an IPv4 IP address for a target iSCSI controller through the source port ID 1. No *tgt_user_name* or *target_chap* value are required.

```
detectiscsistorageportcandidate -srcportid 1 -targetip IP1
```

To view the output specify:

```
lsiscsistorageportcandidate
```

The detailed resulting output:

```
id  src_port_id target_ipv4 target_ipv6 target_iscsiname  iogroup_list configured status site_id site_name
0   1           IP1                IQN1        1:1:1:1        no      Full
```

An invocation example

This example shows target discovery by using an IPv4 IP address for a target iSCSI controller through I/O group 0 and source port ID 0:

```
detectiscsistorageportcandidate -iogrp 0 -srcportid 1 -targetip IP2 -username superuser -chapsecret password2
```

A *tgt_user_name* and *target_chap* value are used. There are only two I/O groups in this system, 0 and 3.

To view the output specify:

```
lsiscsistorageportcandidate
```

The detailed resulting output:

```
id  src_port_id target_ipv4 target_ipv6 target_iscsiname  iogroup_list configured status site_id site_name
0   1           IP2                IQN1        1:-:-:0        no      Full
1   1           IP2                IQN2        1:-:-:0        no      Full
2   1           IP2                IQN3        1:-:-:0        no      Full
```


Isiogrp

Use the **lsiogrp** command to display a concise list or a detailed view of input/output (I/O) groups visible to the clustered system (system).

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 system. Each entry in the list corresponds to a single I/O group.
- The detailed information about a single I/O 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. 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 when using the 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:

```
lsiogrp -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 **lsiogrp** 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.

Description

This command returns a concise list or a detailed view of I/O groups visible to the system.

You can display the following information for this command:

id Indicates the ID of the I/O group.

name Indicates the name of the I/O group.

node_count
Indicates the number of nodes in the I/O group

vdisk_count
Indicates the number of volumes in the I/O group.

host_count
Indicates the number of hosts in the I/O group.

flash_copy_total_memory
Indicates the total amount of memory allocated to FlashCopy.

flash_copy_free_memory
Indicates the total amount of memory allocated to FlashCopy, but unused.

remote_copy_total_memory
Indicates the total amount of memory allocated to Remote Copy, but unused. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap.

remote_copy_free_memory
Indicates the total amount of memory allocated to Remote Copy, but unused. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap.

mirroring_total_memory
Indicates the total amount of memory allocated to mirroring.

raid_total_memory
Indicates the total amount of memory allocated to RAID.

raid_free_memory
Indicates the total amount of memory allocated to RAID, but is unused.

maintenance
Indicates if the I/O group is in maintenance mode. The values are:

- yes
- no

compression_active
Indicates if compression is active for this I/O group.

accessible_vdisk_count
The number of accessible volumes in this I/O group.

compression_supported
Indicates if the I/O group supports compressed volumes.

max_enclosures
Indicates the maximum number of enclosures supported by this I/O group.

encryption_supported
Indicates whether the I/O group supports encryption for attached drives. The possible values are:

- yes
- no

flash_copy_maximum_memory

Identifies the maximum memory that can be set for the specified I/O group. The value must be a numeric string between 552 MB and 2048 MB.

site_id

Identifies the site ID for the I/O group. The possible values are :

- 1
- 2

site_name

Identifies the site name for the I/O group. The value must be an alphanumeric string or blank.

compression_total_memory

Indicates the amount of RAM allocated for compression in the specified I/O group.

fctargetportmode

Indicates the current N_Port ID Virtualization (NPIV) status in the specified I/O group. The values are:

- disabled
- transitional
- enabled
- changing_disabled_to_transitional
- changing_transitional_to_disabled
- changing_enabled_to_transitional
- changing_transitional_to_enabled

A concise invocation example

```
lsiogrp -delim :
```

The resulting output:

```
id:name:node_count:vdisk_count:host_count:site_id:site_name
0:io_grp0:1:0:0:1:chelsea1
1:io_grp1:0:0:0:2:chelsea2
2:io_grp2:0:0:0:3:chelsea1
3:io_grp3:0:0:0:4:chelsea1
4:recovery_io_grp:0:0:0:5:chelsea1
```

A detailed invocation example

```
lsiogrp -delim : 0
```

The detailed 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
raid_total_memory:2MB
raid_free_memory:25.0MB
maintenance: no
compression_active:yes
accessible_vdisk_count:10
compression_supported:yes
max_enclosures:21
```

```
encryption_supported:yes
flash_copy_maximum_memory:2048.0MB
site_id:2
site_name:chelsea2
compression_total_memory:35128.0MB
fctargetportnode:disabled
```

lshbaportcandidate (Deprecated)

The **lshbaportcandidate** command is deprecated. Use either the **lsfcportcandidate** or **lssasportcandidate** command instead.

lsiogrphost

Use the **lsiogrphost** command to display a list of the hosts mapped to a specified I/O group.

Syntax

```
➤ lsiogrphost — [ -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

```
lsiogrphost -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

```
➤— 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

The **lsiogroupcandidate** 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

```
lsiogrpcandidate
```

The resulting output:

```
id
0
1
2
3
4
```

Isiostatsdumps (Deprecated)

Attention: The **lsiostatsdumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

Isiotracedumps (Deprecated)

Attention: The **lsiotracedumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.


Isnode (SVC) / Isnodecanister (Storwize family products)

Use the **lsnode**/ **lsnodecanister** command to return a concise list or a detailed view of nodes or node canisters that are part of the clustered system (system).

The list report style can be used to obtain two styles of report:

- A list containing concise information about all the nodes or node canister on a system. Each entry in the list corresponds to a single node or node canister.
- The detailed information about a single node or node canister.

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 Command-Line Interface (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 (""):


```
lsnode -filtervalue "name=md*"
```

-filtervalue?

Displays a list of valid filter attributes for the **-filtervalue***attribute=value* parameter. The valid filters for the **lsnode** command are:

- canister_id
- config_node/config_nodectanister
- enclosure_id
- enclosure_serial_number
- hardware
- id
- iscsi_alias
- IO_group_name
- IO_group_id
- name
- panel_name
- status
- site_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 object ID or name. 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

This command returns a concise list or a detailed view of nodes or node canisters that are part of the system. Table 33 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 33. Isnode or Isnodecanister attribute values

Attribute	Value
status	Indicates the status. The values are: <ul style="list-style-type: none"> • offline • service • flushing • pending • online • adding • deleting
config_node	Indicates if the node is a configuration node. The values are: <ul style="list-style-type: none"> • yes • no
IO_group_id	Indicates the I/O group of the node.
IO_group_name	Indicates the name of the I/O group in the node.
port_status	Indicates if the node is a configuration node. The values are: <ul style="list-style-type: none"> • active • inactive • not_installed
hardware	Indicates the hardware type (for example, C68).
UPS_serial_number	Indicates the serial number of the UPS.
UPS_unique_id	Indicates the unique ID of the UPS.
panel_name	Indicates the unique identifier for the node.
enclosure_id	Blank. This field is blank for a node-based system.
canister_id	Blank. This field is blank for a node-based system.
enclosure_serial_number	Blank. This field is blank for a node-based system.

Table 33. *Isnode or Isnodecanister attribute values (continued)*

Attribute	Value
service_IP_mode	Indicates the current mode of the service IPv4 <ul style="list-style-type: none"> • Empty if IPv4 is not active • The values are: <ul style="list-style-type: none"> – static (if the service IP is set by the user) – dhcp (if the service IP is set successfully using DHCP server) – dhcpfallback (if the service IP is set to a default value after a DHCP server request failed)
service_IP_mode_6	Indicates the current mode of the service IPv6 <ul style="list-style-type: none"> • Empty if IPv6 is not active • Either static (if the service IP is set by the user) or dhcp (if the service IP set successfully using DHCP server).
site_id	Indicates the site node value.
site_name	Indicates the site name.
identify_LED	Indicates the node or node canister state - on, off, or blank.
product_mtm	Indicates the product machine type.
code_level	Indicates the current level of machine code running on the node. on, off, or blank.
serial_number	Indicates the current serial number for the node.
machine_signature	Indicates the current machine signature for the node.
update_complete	Indicates if the node has completed its update. The values are yes or no.

The first four Fibre Channel (FC) input/output (I/O) ports display the worldwide port name (WWPN), state, and speed. If there are less than four FC I/O ports, the fields display with a WWPN of 0000000000000000, port_status of inactive, and port_speed of N/A. To examine the FC ports, use the **lsportfc** command.

A concise invocation example for SAN Volume Controller

```
lsnode -delim ,
```

The concise resulting output:

```
id,name,UPS_serial_number,WWNN,status,IO_group_id,IO_group_name,config_node,UPS_unique_id,hardware,iscsi_name,iscsi_alias,p
1,node114120,UPS_Fake_SN,5005076801005D00,online,0,io_grp0,yes,10000000000005D00,CG8,iqn.1986-03.com.ibm:2145.mycluster.node
```

A concise invocation example for Storwize V7000

```
lsnodecanister -delim ,
```

The concise resulting output:

```
id,name,UPS_serial_number,WWNN,status,IO_group_id,IO_group_name,config_node,UPS_unique_id,hardware,iscsi_name,iscsi_alias,p
1,node114120,UPS_Fake_SN,5005076801005D00,online,0,io_grp0,yes,10000000000005D00,CG8,iqn.1986-03.com.ibm:2145.mycluster.node
```

A detailed invocation example for SAN Volume Controller

```
lsnode -delim , 1
```

The resulting 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,DH8
iscsi_name,iqn.1986-03.com.ibm:2145.1dcluster-19.h1cn114289
iscsi_alias,
failover_active,no
failover_name,h1cn114253
failover_iscsi_name,iqn.1986-03.com.ibm:2145.1dcluster-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,
service_IP_mode,dhcp
service_IP_mode_6

```

```

site_id,1
site_name,DataCenterA

```

```

identify_LED,on
product_mtm,2145-CG8
code_level,7.4.0.0 (build 99.1.1406102000)
serial_number,78G0123
machine_signature,0123-4567-89AB-CDEF

```

A detailed invocation example for Storwize V7000

```
lsnodecanister -delim , 1
```

The resulting output:

```

id,1
name,h1cn114289
UPS_serial_number,10004BC018
WWNN,5005076801002978
status,online
IO_group_id,0
IO_group_name,io_grp0
partner_nodecanister_id,2
partner_nodecanister_name,h1cn114253
config_nodecanister,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,DH8
iscsi_name,iqn.1986-03.com.ibm:2145.ldcluster-19.hlc114289
iscsi_alias,
failover_active,no
failover_name,hlc114253
failover_iscsi_name,iqn.1986-03.com.ibm:2145.ldcluster-19.hlc114253
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,
service_IP_mode,dhcp
service_IP_mode,
identify_LED,on
site_id,1
site_name,DataCenterA
identify_LED
product_mtm 2145-CG8
code_level,7.4.0.0 (build 99.1.1406102000)
serial_number 78G0123
machine_signature 0123-4567-89AB-CDEF
failover_source
protected_nodes 1,2
```

Isnodebattery

Use the **lsnodebattery** command to display information about the batteries in a node. This command applies to SAN Volume Controller 2145-DH8 systems.

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.

-battery battery_id

(Optional) Specifies the battery ID. If you specify this parameter, you must also specify *node*.

node

(Optional) Specifies the node ID or name.

Description

The command displays information about the batteries in a node. The concise view displays a line for each battery slot in all nodes.

Table 34 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 34. Isnodebattery attribute values

Attribute	Value
node_id	Identifies the node that contains the battery.
node_name	Identifies the node that contains the battery.
battery_id	Identifies the battery in the node.
status	Identifies the status of the battery: <ul style="list-style-type: none">• online if the battery is present and working as usual (this includes a battery calibration).• degraded indicates the battery is present but not working as usual (it has an error logged against it).• offline indicates the battery cannot be detected or is failed (a node error indicating it is missing or failed is logged against the battery). Remember: If the status is offline, all other fields display the most recent battery data. If no data was shown, all fields remain blank.
charging_status	Identifies the charging state of the battery: <ul style="list-style-type: none">• charged indicates the battery is fully charged.• charging indicates the battery is charging.• discharging indicates the battery is losing voltage (life) and it is recalibrating the gas gauge after the battery has recharged..• idle indicates the battery is neither charging nor discharging but it is not fully charged.• reconditioning indicates the battery is reconditioning itself by discharging and subsequently recharging. Important: Reconditioning occurs approximately every three months (on redundant systems) and can take between twelve and forty-eight hours.
recondition_needed	Identifies that the battery needs to be reconditioned or it should start reconditioning soon. A recalibration is required because the reported values from the gas gauge are not trusted. The values are yes and no.
node_percentage_charge	Identifies the battery's percentage charged to determine if it can support the node.

Table 34. *lsnodebattery* attribute values (continued)

Attribute	Value
end_of_life_warning	Identifies the battery's end of life (with a warning noise). The values are yes and no. Important: Replace the battery.
present	Identifies a battery is present. The values are yes and no.
redundant	Identifies if it is safe to remove the battery. The values are yes and no.
remaining_charge_capacity_mAh	Identifies the remaining capacity of the battery (when recharging) in milliamps hours (mAh).
full_charge_capacity_mAh	Identifies the fully-charged capacity of the battery in mAh.
FRU_part_number	Identifies the FRU part number of the battery. The value is contains 7 alphanumeric characters.
FRU_identity	Identifies the 11S number (combining the manufacturing part number and the serial number). The value is contains 22 alphanumeric characters.
compatibility_level	Identifies the battery driver software must support the current software level to operate with this battery (this comes from the battery vital product data or VPD).
last_recondition_timestamp	Indicates a system timestamp for the last successful recalibration of the gas gauge. The format is YYMMDDHHMMSS, where: <ul style="list-style-type: none"> • YY indicates year. • The first MM indicates month. • DD indicates dayr. • HH indicates hour. • The second MM indicates minute. • SS indicates second.
powered_on_hours	Indicates the number of hours the battery has been in a powered node.
cycle_count	Identifies the number of charge or discharge cycles the battery performs.
error_sequence_number	Identifies tthe error log number of the highest priority error. This is generally blank, but if the status is degraded or offline an error sequence number is displayed.

A concise invocation example

If battery 2 in node 1 is not installed:

```
lsnodebattery
```

The resulting output:

node_id	node_name	battery_id	status	charging_status	recondition_needed	node_percentage_charge	end_of_life_warning
1	node1	1	online	charged	no	50	no
1	node1	2	offline				
2	node2	1	online	charged	no	50	no
2	node2	2	online	reconditioning	yes	50	

A concise invocation example

If battery 2 in node 1 is failing to charge:

```
lsnodebattery
```

The resulting output:

node_id	node_name	battery_id	status	charging_status	recondition_needed	node_percentage_charge	end_of_life_warning
1	node1	1	online	charged	no	50	no
1	node1	2	offline	idle	no	50	no
2	node2	1	online	charged	no	100	no
2	node2	2	online	charged	no	100	

A concise invocation example

If battery 2 in node 1 is removed (last known status is presented):

```
lsnodebattery
```

The resulting output:

node_id	node_name	battery_id	status	charging_status	recondition_needed	node_percentage_charge	end_of_life_warning
1	node1	1	online	charged	no	50	no
1	node1	2	offline	charged	no	50	no
2	node2	1	online	charged	no	50	no
2	node2	2	online	reconditioning	yes	50	

A detailed invocation example

```
lsnodebattery -battery 2 2
```

The resulting output:

```
node_id 2
node_name node2
battery_id 2
status online
charging_status reconditioning
present yes
redundant yes
recondition_needed yes
remaining_charge_capacity_mAh 1600
full_charge_capacity_mAh 1950
end_of_life_warning no
FRU_part_number FRU0001
FRU_identity 11SYM30BG123456MAN0001
compatibility_level 5
last_recondition_timestamp 0
powered_on_hours 12345
cycle_count 2
node_percentage_charge 50
error_sequence_number
```

lsnodecandidate (SAN Volume Controller)

Use the **lsnodecandidate** command to list all of the nodes that are available to add to the clustered system.

Syntax

```

>> lsnodecandidate — [ -nohdr ] [ -delim delimiter ] [ -svcconfig ] >>

```

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.

-svcconfig

(Optional) Lists all nodes in the enclosure that are in a candidate state.

Description

Note: The **lsnodecandidate** command is a SAN Volume Controller command. For Storwize V7000, use the **lscontrolenclosurecandidate** command.

This command displays a list of nodes that are available to add to the clustered system. This list includes nodes that are not already part of a clustered system, but are compatible with the clustered system code level. Nodes with hardware types that are incompatible with the installed code are not listed.

The following table describes the possible outputs:

Table 35. lsnodecandidate outputs

Attribute	Description
panel_name	Unique identifier for the node.
UPS_serial_number	The serial number of the UPS.
UPS_unique_id	The unique ID of UPS.
hardware	Describes the type of nodes.
serial_number	Indicates the current serial number for the node (7 characters).
product_mtm	Indicates the current product machine type for the node (8 characters including the hyphen).
machine_signature	Indicates the current machine signature for the node (16-character hexadecimal string with hyphens).

An invocation example

```
lsnodecandidate -delim :
```

The resulting output:

```
id: panel_name:UPS_serial_number:UPS_unique_id:hardware:serial_number:product_mtm:machine_signature
1:146355:10L3ASH:202378101C0D18D8:8G4:78G0123:2145-CG8:0123-4567-89AB-CDEF
```

An invocation example

```
lsnodecandidate
```

The resulting output:

```
id          panel_name  UPS_serial_number  UPS_unique_id  hardware  serial_number  product_mtm  machine_signature
500507680100E85F  168167      UPS_Fake_SN       100000000000E85F  CG8       78G0123       2145-CG8     0123-4567-89AB-CDEF
```

Isnodedependentvdisks (Deprecated)

Attention: The **lsnodedependentvdisks** command is deprecated. Use the **lsdependentvdisks** command instead.

Isnodehw (SVC) / Isnodecanisterhw (Storwize family products)

Use the **Isnodehw** / **Isnodecanisterhw** command to display the configured and actual hardware configuration of nodes in the clustered system.

Syntax

```
►►— Isnodehw — | Isnodecanisterhw — [ -nohdr ] [ -delim — delimiter ] [ object_id — object_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.

object_id | *object_name*

(Required) Specifies the object name or ID.

Description

Table 36 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 36. Attribute values for Isnodehw and Isnodecanisterhw

Attribute	Value
id	Indicates the node or node canister unique ID.
name	Indicates the node or node canister name.
status	Indicates the node or node canister status.
IO_group_id	Indicates the input/output (I/O) group ID.
IO_group_name	Indicates the I/O group name.
hardware	Indicates the hardware model, such as CG8.
actual_different	Indicates if the node or node canister hardware is different from the configured hardware.
actual_valid	Indicates if the node or node canister hardware is valid.
memory_configured	Indicates the configured amount of memory (in GB).
memory_actual	Indicates the currently installed amount of memory (in GB).
memory_valid	Indicates if the actual memory is a valid configuration.
cpu_count	Indicates the maximum number of CPUs for the node.
cpu_socket	Indicates the ID of socket the CPU fields refer to.

Table 36. Attribute values for *lsnodehw* and *lsnodecanisterhw* (continued)

Attribute	Value
cpu_configured	Indicates the configured CPU for this socket.
cpu_actual	Indicates the currently installed CPU in this socket.
cpu_valid	Indicates if the currently installed CPU is a valid configuration.
adapter_count	Indicates the maximum number of adapters for the node (differs by node type).
adapter_location	Indicates the location of this adapter.
adapter_configured	Indicates the configured adapter for this location.
adapter_actual	Indicates the currently installed adapter for this location.
adapter_valid	Indicates if the adapter in this location is valid.
ports_different	Indicates if the current hardware is able to provide more I/O ports? The values are yes and no.

An invocation example for SAN Volume Controller

```
lsnodehw -delim , 1
```

The resulting output:

```
id,1
name,h1cn114289
status,online
IO_group_id,0
IO_group_name,io_grp0
hardware,CG8
actual_different,yes
actual_valid,no
memory_configured,8
memory_actual,8
memory_valid,yes
cpu_count,2
cpu_socket,1
cpu_configured,4 core Intel(R) Xeon(R) CPU E3110 @ 3.0GHz
cpu_actual,4 core Intel(R) Xeon(R) CPU E3110 @ 3.0GHz
cpu_valid,yes
cpu_socket,2
cpu_configured,none
cpu_actual,none
cpu_valid,yes
adapter_count,4
adapter_location,0
adapter_configured,1Gb/s Ethernet adapter
adapter_actual,1Gb/s Ethernet adapter
adapter_valid,yes
adapter_location,0
adapter_configured,1Gb/s Ethernet adapter
adapter_actual,1Gb/s Ethernet adapter
adapter_valid,yes
adapter_location,1
adapter_configured,Four port 8Gb/s FC adapter card
adapter_actual,Four port 8Gb/s FC adapter card
adapter_valid,yes
adapter_location,2
adapter_configured,none
adapter_actual,Four port 8Gb/s FC adapter card
adapter_valid,no
ports_different yes
```


An invocation example for Storwize V7000

```
lsnodecanisterhw -delim , 1
```

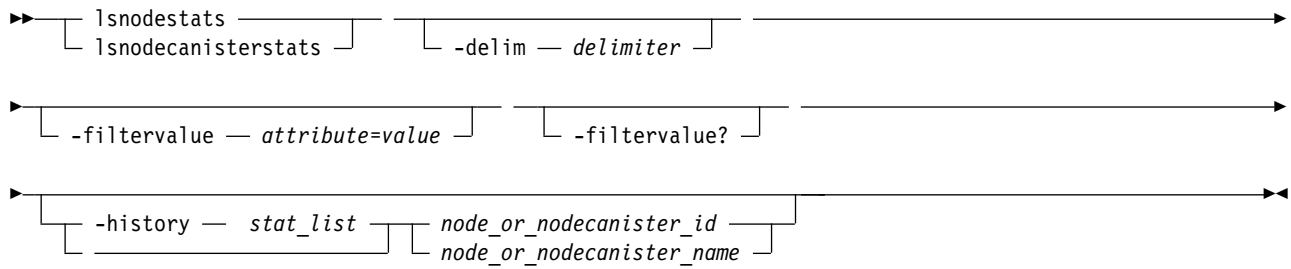
The resulting output

```
id,1
name,h1cn114289
status,online
IO_group_id,0
IO_group_name,io_grp0
hardware,112
...
```

Isnodestats (SVC) / Isnodecanisterstats (Storwize family products)

Use the **lsnodestats** / **lsnodecanisterstats** command to display the most recent values of statistics for all nodes or node canisters, and display all statistics for a particular node or node canister. Additionally, You can use this command to display a history of values for a given subset of available statistics.

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.

-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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):

```
lsenclosurestats -filtervalue stat_name=temp_f
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue attribute=value** parameter:

- node_id
- node_name

- stat_name

-history stat_list

(Optional) Provides a table of statistical values for the specified node. The *stat_list* is a colon-delimited list of one or more statistical values. A table is generated for each entry in the *stat_list*.

Remember: If **-history** is specified, a node ID or name must be specified as well.

node_or_nodecanister_id | node_or_nodecanister_name

(Optional) Identifies the node or node canister for which you want to request statistics.

Description

This command returns a concise list or a detailed view of nodes or node canisters that are part of the clustered system. Table 37 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 37. Attribute values for lsnodestats or lsnodecanister

Attribute	Value
node_id	The ID of the node or node canister.
node_name	The name of the node or node canister.
stat_current	The current value of the statistic field.
stat_list	The system history of the reported statistics. The list of statistics can contain multiple items separated by colons.
stat_name	The name of the statistic field. See Table 38 on page 226 for descriptions of available statistics.
stat_peak	The peak value of the statistic field in the last five minutes.
stat_peak_time	The time that the peak occurred.
sample_time	The time of the sample occurrence.
stat_value	The statistical value at the epoch interval.

Remember: Filtering is supported on the stat_name field using the concise view.

An invocation example

lsnodestats

The resulting output:

```
node_id node_name stat_name      stat_current stat_peak stat_peak_time
1       node1     cpu_pc          5            9         111123105330
1       node1     fc_mb          218          238         111123105440
1       node1     fc_io          1122         1501         111123105435
1       node1     sas_mb         282          402         111123105335
1       node1     sas_io         3129         4427         111123105335
1       node1     iscsi_mb       0            0         111123105825
1       node1     iscsi_io       0            0         111123105825
1       node1     write_cache_pc 0            0         111123105825
1       node1     total_cache_pc 0            0         111123105825
1       node1     vdisk_mb       218          238         111123105440
1       node1     vdisk_io       1076         1452         111123105435
1       node1     vdisk_ms       52           60         111123105605
1       node1     mdisk_mb       218          238         111123105435
1       node1     mdisk_io       1874         2386         111123105435
1       node1     mdisk_ms       15           33         111123105605
1       node1     drive_mb       281          401         111123105335
1       node1     drive_io       3130         4060         111123105335
```

1	node1	drive_ms	13	27	111123105605
1	node1	vdisk_r_mb	134	157	111123105440
1	node1	vdisk_r_io	561	885	111123105430
1	node1	vdisk_r_ms	37	45	111123105605
1	node1	vdisk_w_mb	84	89	111123105700
1	node1	vdisk_w_io	515	587	111123105625
1	node1	vdisk_w_ms	67	84	111123105330
1	node1	mdisk_r_mb	133	155	111123105510
1	node1	mdisk_r_io	1337	1789	111123105435
1	node1	mdisk_r_ms	15	33	111123105605
1	node1	mdisk_w_mb	84	89	111123105700
1	node1	mdisk_w_io	536	611	111123105625
1	node1	mdisk_w_ms	17	32	111123105605
1	node1	drive_r_mb	151	295	111123105335
1	node1	drive_r_io	1700	2904	111123105335
1	node1	drive_r_ms	14	30	111123105605
1	node1	drive_w_mb	130	137	111123105700
1	node1	drive_w_io	1429	1586	111123105625
1	node1	drive_w_ms	12	22	111123105605
1	node1	iplink_mb	0	1	130523104536
1	node1	iplink_io	0	10	130523104536
2	node2	cpu_pc	6	7	111123105624
2	node2	fc_mb	132	145	111123105724
2	node2	fc_io	1519	1944	111123105739
2	node2	sas_mb	189	308	111123105619
2	node2	sas_io	2737	4099	111123105614
2	node2	iscsi_mb	0	0	111123105824
2	node2	iscsi_io	0	0	111123105824
2	node2	write_cache_pc	0	0	111123105824
2	node2	total_cache_pc	0	0	111123105824
2	node2	vdisk_mb	132	145	111123105724
2	node2	vdisk_io	1459	1892	111123105739
2	node2	vdisk_ms	47	81	111123105514
2	node2	mdisk_mb	132	145	111123105724
2	node2	mdisk_io	1635	2066	111123105739
2	node2	mdisk_ms	8	18	111123105619
2	node2	drive_mb	189	310	111123105619
2	node2	drive_io	2735	3750	111123105619
2	node2	drive_ms	9	20	111123105604
2	node2	vdisk_r_mb	20	21	111123105809
2	node2	vdisk_r_io	796	1180	111123105739
2	node2	vdisk_r_ms	2	8	111123105529
2	node2	vdisk_w_mb	112	134	111123105349
2	node2	vdisk_w_io	662	805	111123105504
2	node2	vdisk_w_ms	100	104	111123105624
2	node2	mdisk_r_mb	20	21	111123105809
2	node2	mdisk_r_io	951	1330	111123105739
2	node2	mdisk_r_ms	2	7	111123105529
2	node2	mdisk_w_mb	112	134	111123105349
2	node2	mdisk_w_io	684	834	111123105504
2	node2	mdisk_w_ms	16	36	111123105619
2	node2	drive_r_mb	17	132	111123105619
2	node2	drive_r_io	899	1920	111123105619
2	node2	drive_r_ms	6	12	111123105344
2	node2	drive_w_mb	171	206	111123105504
2	node2	drive_w_io	1837	2230	111123105504
2	node2	drive_w_ms	11	26	111123105619
1	node1	iplink_mb	0	1	130523104536
1	node1	iplink_io	0	10	130523104536
cloud_up_mb	0	0	161118051715		
cloud_up_ms	0	0	161118051715		
cloud_down_mb	0	0	161118051715		
cloud_down_ms	0	0	161118051715		

An invocation example of a node-based, filtered invocation example for lsnodestats

```
lsnodestats -filtervalue stat_name=sas_io:stat_name=sas_mb node1
```

The resulting output:

node_id	node_name	stat_name	stat_current	stat_peak	stat_peak_time
1	node1	sas_mb	212	421	111123105840
1	node1	sas_io	2477	4184	111123105840

An invocation example of an historical view that can list multiple statistics and requires a node-based invocation

```
lsnodestats -history cpu_pc:fc_mb:sas_mb node1
```

The resulting output:

node_id	node_name	sample_time	stat_name	stat_value
2	node2	111123105839	cpu_pc	6
2	node2	111123105844	cpu_pc	5
2	node2	111123105849	cpu_pc	5
2	node2	111123105854	cpu_pc	5
2	node2	111123105859	cpu_pc	6
2	node2	111123105904	cpu_pc	5
2	node2	111123105909	cpu_pc	5
2	node2	111123105914	cpu_pc	5
2	node2	111123105919	cpu_pc	5
2	node2	111123105924	cpu_pc	5
2	node2	111123105929	cpu_pc	5
2	node2	111123105934	cpu_pc	5
2	node2	111123105839	fc_mb	128
2	node2	111123105844	fc_mb	126
2	node2	111123105849	fc_mb	123
2	node2	111123105854	fc_mb	142
2	node2	111123105859	fc_mb	119
2	node2	111123105904	fc_mb	131
2	node2	111123105909	fc_mb	157
2	node2	111123105914	fc_mb	177
2	node2	111123105919	fc_mb	182
2	node2	111123105924	fc_mb	182
2	node2	111123105929	fc_mb	155
2	node2	111123105934	fc_mb	177
2	node2	111123105839	sas_mb	191
2	node2	111123105844	sas_mb	191
2	node2	111123105849	sas_mb	185
2	node2	111123105854	sas_mb	216
2	node2	111123105859	sas_mb	181
2	node2	111123105904	sas_mb	198
2	node2	111123105909	sas_mb	228
2	node2	111123105914	sas_mb	243
2	node2	111123105919	sas_mb	251
2	node2	111123105924	sas_mb	248
2	node2	111123105929	sas_mb	217
2	node2	111123105934	sas_mb	242

The following table provides the possible values that are applicable to the values that are displayed for the **stat_name** attribute.

Table 38. Stat_name field values

Value	Description
compression_cpu_pc	Displays the percentage of allocated CPU capacity utilized for compression.
cpu_pc	Displays the percentage of allocated CPU capacity utilized for the system.

Table 38. Stat_name field values (continued)

Value	Description
fc_mb	Displays the total number of megabytes transferred per second for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.
fc_io	Displays the total input/output (I/O) operations transferred per seconds for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.
sas_mb	Displays the total number of megabytes transferred per second for serial-attached SCSI (SAS) traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
sas_io	Displays the total I/O operations transferred per second for SAS traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
iscsi_mb	Displays the total number of megabytes transferred per second for iSCSI traffic on the system.
iscsi_io	Displays the total I/O operations transferred per second for iSCSI traffic on the system.
write_cache_pc	Displays the percentage of the write cache usage for the node.
total_cache_pc	Displays the total percentage for both the write and read cache usage for the node.
vdisk_mb	Displays the average number of megabytes transferred per second for read and write operations to volumes during the sample period.
vdisk_io	Displays the average amount of I/O operations transferred per second for read and write operations to volumes during the sample period.
vdisk_ms	Displays the average amount of time in milliseconds that the system takes to respond to read and write requests to volumes over the sample period.
mdisk_mb	Displays the average number of megabytes transferred per second for read and write operations to MDisks during the sample period.
mdisk_io	Displays the average amount of I/O operations transferred per second for read and write operations to MDisks during the sample period.
mdisk_ms	Displays the average amount of time in milliseconds that the system takes to respond to read and write requests to MDisks over the sample period.
drive_mb	Displays the average number of megabytes transferred per second for read and write operations to drives during the sample period
drive_io	Displays the average amount of I/O operations transferred per second for read and write operations to drives during the sample period.
drive_ms	Displays the average amount of time in milliseconds that the system takes to respond to read and write requests to drives over the sample period.
vdisk_w_mb	Displays the average number of megabytes transferred per second for read and write operations to volumes during the sample period.
vdisk_w_io	Displays the average amount of I/O operations transferred per second for write operations to volumes during the sample period.
vdisk_w_ms	Displays the average amount of time in milliseconds that the system takes to respond to write requests to volumes over the sample period.
mdisk_w_mb	Displays the average number of megabytes transferred per second for write operations to MDisks during the sample period.
mdisk_w_io	Displays the average amount of I/O operations transferred per second for write operations to MDisks during the sample period.
mdisk_w_ms	Displays the average amount of time in milliseconds that the system takes to respond to write requests to MDisks over the sample period.

Table 38. Stat_name field values (continued)

Value	Description
drive_w_mb	Displays the average number of megabytes transferred per second for write operations to drives during the sample period
drive_w_io	Displays the average amount of I/O operations transferred per second for write operations to drives during the sample period.
drive_w_ms	Displays the average amount of time in milliseconds that the system takes to respond to write requests to drives over the sample period.
vdisk_r_mb	Displays the average number of megabytes transferred per second for read operations to volumes during the sample period.
vdisk_r_io	Displays the average amount of I/O operations transferred per second for read operations to volumes during the sample period.
vdisk_r_ms	Displays the average amount of time in milliseconds that the system takes to respond to read requests to volumes over the sample period.
mdisk_r_mb	Displays the average number of megabytes transferred per second for read operations to MDisk during the sample period.
mdisk_r_io	Displays the average amount of I/O operations transferred per second for read operations to MDisk during the sample period.
mdisk_r_ms	Displays the average amount of time in milliseconds that the system takes to respond to read requests to MDisk over the sample period.
drive_r_mb	Displays the average number of megabytes transferred per second for read operations to drives during the sample period
drive_r_io	Displays the average amount of I/O operations transferred per second for read operations to drives during the sample period.
drive_r_ms	Displays the average amount of time in milliseconds that the system takes to respond to read requests to drives over the sample period.
iplink_mb	The total number of megabytes transferred per second for Internet Protocol (IP) replication traffic on the system. This value does not include iSCSI host input/output (I/O) operations.
iplink_comp_mb	Displays the average number of compressed megabytes transferred per second (MBps) over the IP replication link during the sample period.
iplink_io	The total input/output (I/O) operations transferred per second for IP partnership traffic on the system. This value does not include Internet Small Computer System Interface (iSCSI) host I/O operations.
cloud_up_mb	Displays the average number of megabits transferred per second (Mbps) for upload operations to a cloud account during the sample period.
cloud_up_ms	Displays the average amount of time (in milliseconds) it takes for the system to respond to upload requests to a cloud account during the sample period.
cloud_down_mb	Displays the average number of Mbps for download operations to a cloud account during the sample period.
cloud_down_ms	Displays the average amount of time (in milliseconds) it takes for the system to respond to download requests to a cloud account during the sample period.

An invocation example

```
lsnodecanisterstats
```

The resulting output:

node_id	node_name	stat_name	stat_current	stat_peak	stat_peak_time
1	node1	cpu_pc	5	9	111123105330
1	node1	fc_mb	218	238	111123105440

1	node1	fc_io	1122	1501	111123105435
1	node1	sas_mb	282	402	111123105335
1	node1	sas_io	3129	4427	111123105335
1	node1	iscsi_mb	0	0	111123105825
1	node1	iscsi_io	0	0	111123105825
1	node1	write_cache_pc	0	0	111123105825
1	node1	total_cache_pc	0	0	111123105825
1	node1	vdisk_mb	218	238	111123105440
1	node1	vdisk_io	1076	1452	111123105435
1	node1	vdisk_ms	52	60	111123105605
1	node1	mdisk_mb	218	238	111123105435
1	node1	mdisk_io	1874	2386	111123105435
1	node1	mdisk_ms	15	33	111123105605
1	node1	drive_mb	281	401	111123105335
1	node1	drive_io	3130	4060	111123105335
1	node1	drive_ms	13	27	111123105605
1	node1	vdisk_r_mb	134	157	111123105440
1	node1	vdisk_r_io	561	885	111123105430
1	node1	vdisk_r_ms	37	45	111123105605
1	node1	vdisk_w_mb	84	89	111123105700
1	node1	vdisk_w_io	515	587	111123105625
1	node1	vdisk_w_ms	67	84	111123105330
1	node1	mdisk_r_mb	133	155	111123105510
1	node1	mdisk_r_io	1337	1789	111123105435
1	node1	mdisk_r_ms	15	33	111123105605
1	node1	mdisk_w_mb	84	89	111123105700
1	node1	mdisk_w_io	536	611	111123105625
1	node1	mdisk_w_ms	17	32	111123105605
1	node1	drive_r_mb	151	295	111123105335
1	node1	drive_r_io	1700	2904	111123105335
1	node1	drive_r_ms	14	30	111123105605
1	node1	drive_w_mb	130	137	111123105700
1	node1	drive_w_io	1429	1586	111123105625
1	node1	drive_w_ms	12	22	111123105605
1	node1	iplink_mb	0	1	130523104536
1	node1	iplink_io	0	10	130523104536
2	node2	cpu_pc	6	7	111123105624
2	node2	fc_mb	132	145	111123105724
2	node2	fc_io	1519	1944	111123105739
2	node2	sas_mb	189	308	111123105619
2	node2	sas_io	2737	4099	111123105614
2	node2	iscsi_mb	0	0	111123105824
2	node2	iscsi_io	0	0	111123105824
2	node2	write_cache_pc	0	0	111123105824
2	node2	total_cache_pc	0	0	111123105824
2	node2	vdisk_mb	132	145	111123105724
2	node2	vdisk_io	1459	1892	111123105739
2	node2	vdisk_ms	47	81	111123105514
2	node2	mdisk_mb	132	145	111123105724
2	node2	mdisk_io	1635	2066	111123105739
2	node2	mdisk_ms	8	18	111123105619
2	node2	drive_mb	189	310	111123105619
2	node2	drive_io	2735	3750	111123105619
2	node2	drive_ms	9	20	111123105604
2	node2	vdisk_r_mb	20	21	111123105809
2	node2	vdisk_r_io	796	1180	111123105739
2	node2	vdisk_r_ms	2	8	111123105529
2	node2	vdisk_w_mb	112	134	111123105349
2	node2	vdisk_w_io	662	805	111123105504
2	node2	vdisk_w_ms	100	104	111123105624
2	node2	mdisk_r_mb	20	21	111123105809
2	node2	mdisk_r_io	951	1330	111123105739
2	node2	mdisk_r_ms	2	7	111123105529
2	node2	mdisk_w_mb	112	134	111123105349
2	node2	mdisk_w_io	684	834	111123105504
2	node2	mdisk_w_ms	16	36	111123105619
2	node2	drive_r_mb	17	132	111123105619

2	node2	drive_r_io	899	1920	111123105619
2	node2	drive_r_ms	6	12	111123105344
2	node2	drive_w_mb	171	206	111123105504
2	node2	drive_w_io	1837	2230	111123105504
2	node2	drive_w_ms	11	26	111123105619
1	node1	iplink_mb	0	1	130523104536
1	node1	iplink_io	0	10	130523104536
iplink_comp_mb		0	250	151014133723	

An invocation example of a node-based, filtered invocation example for lsnodecanisterstats

```
lsnodecanisterstats -filtervalue stat_name=sas_io:stat_name=sas_mb node1
```

The resulting output:

node_id	node_name	stat_name	stat_current	stat_peak	stat_peak_time
1	node1	sas_mb	212	421	111123105840
1	node1	sas_io	2477	4184	111123105840

An invocation example of an historical view that can list multiple statistics and requires a node-based invocation

```
lsnodecanisterstats -history cpu_pc:fc_mb:sas_mb node1
```

The resulting output:

node_id	node_name	sample_time	stat_name	stat_value
2	node2	111123105839	cpu_pc	6
2	node2	111123105844	cpu_pc	5
2	node2	111123105849	cpu_pc	5
2	node2	111123105854	cpu_pc	5
2	node2	111123105859	cpu_pc	6
2	node2	111123105904	cpu_pc	5
2	node2	111123105909	cpu_pc	5
2	node2	111123105914	cpu_pc	5
2	node2	111123105919	cpu_pc	5
2	node2	111123105924	cpu_pc	5
2	node2	111123105929	cpu_pc	5
2	node2	111123105934	cpu_pc	5
2	node2	111123105839	fc_mb	128
2	node2	111123105844	fc_mb	126
2	node2	111123105849	fc_mb	123
2	node2	111123105854	fc_mb	142
2	node2	111123105859	fc_mb	119
2	node2	111123105904	fc_mb	131
2	node2	111123105909	fc_mb	157
2	node2	111123105914	fc_mb	177
2	node2	111123105919	fc_mb	182
2	node2	111123105924	fc_mb	182
2	node2	111123105929	fc_mb	155
2	node2	111123105934	fc_mb	177
2	node2	111123105839	sas_mb	191
2	node2	111123105844	sas_mb	191
2	node2	111123105849	sas_mb	185
2	node2	111123105854	sas_mb	216
2	node2	111123105859	sas_mb	181
2	node2	111123105904	sas_mb	198
2	node2	111123105909	sas_mb	228
2	node2	111123105914	sas_mb	243
2	node2	111123105919	sas_mb	251
2	node2	111123105924	sas_mb	248
2	node2	111123105929	sas_mb	217
2	node2	111123105934	sas_mb	242

lsnodevpd (SVC) / lsnodecanistervpd (Storwize family products)

Use the **lsnodevpd** / **lsnodecanistervpd** command to display the vital product data (VPD) for each node.

Syntax

```
▶▶— lsnodevpd — | lsnodecanistervpd — ————┐ ————┐—————▶
                                     └ -nohdr ┘   └ -delim — delimiter ┘

▶┐ object_id —————▶
  └ object_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.

object_id | *object_name*

(Required) Specifies the object name or ID.

Description

This command displays the VPD for the specified node or node canister. 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 SAN Volume Controller 2145-8G4 nodes, the VPD displays the device serial number of the Fibre Channel adapter as N/A.

Table 39. Attribute values for `lsnodevpd` and `lsnodecanistervpd`

Value	Description
system board	Displays the system board information.
part_number	Displays the total number of system part numbers.
system_serial_number	Displays the total number of system serial numbers.
number_of_processors	Displays the total number of system processors.
number_of_memory_modules	Displays the total number of memory modules.
number_of_fans	Displays the total number of system fans.
number_of_FC_cards	Displays the total number of Fibre Channel (FC) cards.
number_of_Ethernet_cards	Displays the total number of Ethernet cards.

An invocation example for SAN Volume Controller

```
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_generic_devices 3
number_of_FC_adapters 1
number_of_Ethernet_adapters 3
number_of_SAS_adapters 0
number_of_Bus_adapters 0
number_of_power_supplies 2
number_of_local_managed_disks 0
BIOS_manufacturer IBM Corp.
BIOS_version -[D6E124AUS-1.01]-
BIOS_release_date 04/30/2009
system_manufacturer IBM
system_product System x3650 M4 -[2145DH8]-
version 00
system_product IBM System x -[2145CF8]-
planar_manufacturer IBM
CMOS_battery_part_number 33F8354
frame_assembly_part_number
power_cable_assembly_part_number 31P1294
service_processor_firmware 1.01
disk_controller 44E8690

processor: 6 fields
part_number 46D1266
processor_location Processor 1
```

```

manufacturer Intel(R) Corporation
version Intel(R) Xeon(R) CPU           E5530  @ 2.40GHz
speed 2400
status Enabled
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

part_number 44T1493
device_location DIMM04
bank_location BANK04
size (MB) No Module Installed
manufacturer Not Specified
serial_number Not Specified

part_number 44T1493
device_location DIMM05
bank_location BANK05
size (MB) 4096
manufacturer Hynix
serial_number 12F41112

part_number 44T1493
device_location DIMM06
bank_location BANK06
size (MB) 4096
manufacturer Hynix
serial_number 2AF41112

part_number 44T1493
device_location DIMM07
bank_location BANK07
size (MB) 4096
manufacturer Hynix
serial_number D128312E

part_number 44T1493
device_location DIMM08
bank_location BANK08
size (MB) 4096
manufacturer Hynix
serial_number D028C12E

part_number 44T1493
device_location DIMM09
bank_location BANK09
size (MB) No Module Installed
manufacturer Not Specified
serial_number Not Specified

```

part_number 44T1493
device_location DIMM10
bank_location BANK10
size (MB) No Module Installed
manufacturer Not Specified
serial_number Not Specified

part_number 44T1493
device_location DIMM11
bank_location BANK11
size (MB) No Module Installed
manufacturer Not Specified
serial_number Not Specified

part_number 44T1493
device_location DIMM12
bank_location BANK12
size (MB) No Module Installed
manufacturer Not Specified
serial_number Not Specified

part_number 44T1493
device_location DIMM13
bank_location BANK13
size (MB) No Module Installed
manufacturer Not Specified
serial_number Not Specified

part_number 44T1493
device_location DIMM14
bank_location BANK14
size (MB) No Module Installed
manufacturer Not Specified
serial_number Not Specified

part_number 44T1493
device_location DIMM15
bank_location BANK15
size (MB) No Module Installed
manufacturer Not Specified
serial_number Not Specified

part_number 44T1493
device_location DIMM16
bank_location BANK16
size (MB) No Module Installed
manufacturer Not Specified
serial_number Not Specified

fan: 12 fields
part_number 43V6929
location location1

part_number 43V6929
location location2

part_number 43V6929
location location3

part_number 43V6929
location location4

part_number 43V6929
location location5

part_number 43V6929

location location6

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

Fibre channel port: 44 fields

part_number 31P1338
manufacturer JDSU
device PLRXPLVCSH423N
serial_number C945VK0RB
supported_speeds 2,4,8 Gbps
connector_type LC
transmitter_type SN
wavelength 850
max_distance_by_cable_type OM1:20,OM2:50,OM3:150
hw_revision 2
port_number 1

part_number 31P1338
manufacturer JDSU
device PLRXPLVCSH423N
serial_number C945VK0KU
supported_speeds 2,4,8 Gbps
connector_type LC
transmitter_type SN
wavelength 850
max_distance_by_cable_type OM1:20,OM2:50,OM3:150
hw_revision 2
port_number 2

part_number 31P1338
manufacturer JDSU
device PLRXPLVCSH423N
serial_number C945VK0KT
supported_speeds 2,4,8 Gbps
connector_type LC
transmitter_type SN
wavelength 850
max_distance_by_cable_type OM1:20,OM2:50,OM3:150
hw_revision 2
port_number 3

part_number 31P1338
manufacturer JDSU
device PLRXPLVCSH423N
serial_number C945VK0RA
supported_speeds 2,4,8 Gbps
connector_type LC
transmitter_type SN
wavelength 850
max_distance_by_cable_type OM1:20,OM2:50,OM3:150
hw_revision 2
port_number 4

Adapter card: 9 fields

card_type Ethernet
part_number 43V7072
port_numbers 1 2
location 0

device_serial_number 0123456789
manufacturer Unknown
device NetXtreme II Gigabit Ethernet
card_revision Unknown
chip_revision 8.0

Ethernet port: 22 fields
part_number Unknown
manufacturer N/A
device N/A
serial_number N/A
supported_speeds 10,100 Mbps,1 Gbps
connector_type N/A
transmitter_type N/A
wavelength N/A
max_distance_by_cable_type N/A
hw_revision N/A
port_number 1

part_number Unknown
manufacturer N/A
device N/A
serial_number N/A
supported_speeds 10,100 Mbps,1 Gbps
connector_type N/A
transmitter_type N/A
wavelength N/A
max_distance_by_cable_type N/A
hw_revision N/A
port_number 2

Adapter card: 9 fields
card_type Ethernet
part_number 31P1559
port_numbers 3 4
location 2
device_serial_number BT05149496
manufacturer Emulex Corp
device Emulex/OneConnect 10Gb NIC (be3)
card_revision 1.0
chip_revision 0.2

Ethernet port: 22 fields
part_number 31P1549
manufacturer FINISAR CORP.
device FTLX8571D3BCL
serial_number AHE05K7
supported_speeds 10 Gbps
connector_type LC
transmitter_type 10G Base-SR
wavelength 850
max_distance_by_cable_type OM1:30,OM2:80,OM3:300
hw_revision A
port_number 3

part_number 31P1549
manufacturer JDSU
device PLRXPLSCS4321N
serial_number C825UB0D2
supported_speeds 10 Gbps
connector_type LC
transmitter_type 10G Base-SR
wavelength 850
max_distance_by_cable_type OM1:30,OM2:80,OM3:300
hw_revision 1
port_number 4

```

device: 24 fields
part_number 31P1339
bus USB
device 0
model IBM USB Endeavour
revision 1.1
serial_number NA
approx_capacity 0
hw_revision 0

part_number 42D0673
bus scsi
device 0
model MBE2073RC
revision SC13
serial_number D3A01C0HSC13SC13SC1
approx_capacity 68
hw_revision

part_number N/A
bus scsi
device 0
model STEC USB 2.0
revision 1113
serial_number NA
approx_capacity 1
hw_revision

system code level: 4 fields
id 58
node_name dvt151769
WWNN 0x500507680100b7d2
code_level 6.4.1.3 (build 75.0.1212193000)
object_name_model

front panel assembly: 3 fields
front_panel_id 151769
part_number N/A
...
battery_midplane_FRU_part 12Z9880
battery_midplane_part_identity 11S98Z1230YM11RM234567
battery_midplane_FW_version 1.6
battery_power_cable_FRU_part 12Z9881
battery_power_sense_cable_FRU_part 12Z9882
battery_comms_cable_FRU_part 12Z9883
battery_EP0W_cable_FRU_part 12Z9884

```

An invocation example for Storwize V7000

```
lsnodecanistervpd 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_Ethernet_cards 3
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 System x3650 M4 -[2145DH8]-
version 00planar_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)
nodecanister_name nodecanister1
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
front_panel_id 161040
front_panel_locale en_US

part_number N/A

UPS: 10 fields
electronics_assembly_part_number 64P8326
battery_part_number 31P0710
battery: 7 fields
battery_midplane_FRU_part 12Z9880
battery_midplane_part_identity 11S98Z1230YM11RM234567
battery_midplane_FW_version 1.6
battery_power_cable_FRU_part 12Z9881

```

```

battery_power_sense_cable_FRU_part 12Z9882
battery_comms_cable_FRU_part 12Z9883
battery_EPOW_cable_FRU_part 12Z9884
UPS_assembly_part_number 64P8326
input_power_cable_part_number CountryDependent
UPS_serial_number 1000840050
UPS_type 2145UPS 1U
UPS_internal_part_number P31P0875
UPS_unique_id 0x20400002047c0140
UPS_main_firmware 1.02
UPS_comms_firmware 1.20
...

```

Isportusb

Use the **lspportusb** command to display information about Universal Serial Bus (USB) ports.

Syntax

```

▶▶ lspportusb — [ -nohdr ] [ -delim delimiter ] [ usb_port_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.

usb_port_id

(Optional) Specifies the USB port ID. Used when detailed information about a USB port is required.

Description

This command displays information about Universal Serial Bus (USB) ports.

Table 40 provides the attribute values that can be displayed as output view data.

Table 40. *lspportusb* output

Attribute	Possible Values
id	Indicates the unique ID of the USB port in the system. This is the <code>usb_port_id</code> shown in the syntax above. The value is a numeric 0 or greater.
node_id	Indicates the ID of the node where the USB port is. The value is a numeric string.

Table 40. *Isportusb* output (continued)

Attribute	Possible Values
node_name	Indicates the name of the node where the USB port is. The value is an alphanumeric string.
node_side	Indicates the side of the node where the USB port is. The values are front and rear.
port_id	Indicates the ID of the USB port on the node side. The value is a numeric 1 or greater.
status	Indicates the status of the USB port. The values are: <ul style="list-style-type: none"> • active, which indicates that a USB flash drive is plugged in and can be used by the system. • inactive, which indicates that no USB flash drive is detected. • unsupported, which indicates that a USB device is plugged in but cannot be used.
encryption_state	Indicates the encryption status of the USB device attached to the port. The values are: <ul style="list-style-type: none"> • Blank indicates that it is not in use for encryption • validated indicates that encryption keys are present and validated • missing indicates that encryption keys were validated and were then removed, and the DMP has not yet run to confirm the absence. • prepared indicates that encryption keys are prepared as part of a rekey operation. • validated_prepared indicates that encryption keys have been validated and prepared as part of a rekey operation. • wrong_system indicates that encryption keys are detected on the USB device but none valid for the system. • old indicates that the USB device contains encryption keys that were generated for this system - but they are not the current keys. • error indicates that an encryption key is detected and something might be wrong with it.
encryption_filename	Indicates the name of the file in the rot directory of the USB device to which the encryption state relates. The file name can contain up to 110 characters.
service_state	Indicates the USB command status. The values are: <ul style="list-style-type: none"> • Blank indicates that no command is active. • running indicates that satask.txt is processing, and default USB processing is ongoing. • complete indicates that satask.txt is processing and default USB processing is complete. • install_image indicates that satask.txt processing has not started because there is an installation image on the USB flash drive.

A concise invocation example

```
lsportusb
```

The resulting output:

```
id:node_id:node_name:node_side:port_id:status:service_state
0:1:node1:rear:1:inactive
1:1:node1:rear:2:active:validated:complete
2:2:node2:rear:1:active::complete
3:2:node2:rear:2:active:wrong_system:complete
```

A detailed invocation example

```
lsportusb 3
```

The resulting output:

```
id 3
node_id 2
node_name node2
node_side rear
port_id 2
status active

encryption_state wrong_system
encryption_filename encryption_key_filename_BadSystem
service_state complete
```

lsportip

Use the **lsportip** command to list the Internet Small Computer Systems Interface (iSCSI) Internet Protocol (IP) addresses assigned for each port on each node in the clustered system.

Syntax

```
➤— lsportip — [ -filtervalue — attribute=value ] [ -filtervalue? ] [ -nohdr ] ➤
➤ [ -delim — delimiter ] [ ethernet_port_id ] ➤
```

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, enclose the filter entry within double quotation marks (""), as follows:

```
lsportip -filtervalue "node_name=md*"
```

-filtervalue?

(Optional) Displays the valid filter attributes. The following filter attributes for the **lsportip** command are valid:

- id

- `node_id`
- `node_name`
- `state`
- `failover`

-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 (1, 2, 3 or 4). 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 clustered system. The concise view displays two rows of output for each Ethernet port. Each node has two Ethernet ports.

Use the **lsportip** 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 Internet Small Computer System Interface, an Internet Protocol (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.

This command enables you to view information about system port status.

Table 41 on page 244 provides the attribute values that can be displayed as output view data.

Table 41. *Isportip* output

Attribute	Description
id	Indicates the ID of the Ethernet port.
node_id	Indicates the ID of the node that contains the port.
node_name	Indicates the name of the node that contains the port.
IP address	Indicates the IPv4 address (and is blank if there is none).
mask	Indicates the IPv4 subnet mask (and is blank if there is none).
gateway	Indicates the IPv4 gateway (and is blank if there is none).
IP_address_6	Indicates the IPv6 address (and is blank if there is none).
prefix_6	Indicates the IPv6 prefix (and is blank if there is none).
gateway_6	Indicates the IPv6 gateway address (and is blank if there is none).
MAC	Indicates the current MAC address (blank if unknown).
duplex	Indicates the current duplex state of the port (blank if unknown).
state	Indicates the state of iSCSI addresses. The values can be: <ul style="list-style-type: none"> • unconfigured: There is no iSCSI address (or hardware might not exist). • configured: The iSCSI address is configured. • management_only: It is not configurable for I/O operations.
link_state	Indicates the link state of ethernet port. The values are active and inactive.
host	Displays the IPv4 address used for host attach.
remote_copy	Displays the IPv4 remote copy port group ID. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap.
host_6	Displays the IPv6 address used for host attach.
remote_copy_6	Displays the IPv6 remote copy port group ID. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap.
remote_copy_status	Displays the IPv4 remote copy status. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap.
remote_copy_status_6	Displays the IPv6 remote copy status. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap.
vlan	Displays the virtual local area network (VLAN) ID associated with the IPv4 address on this port (a numeric character from 1 to 4094).
vlan_6	Displays the VLAN ID associated with the IPv6 address on this port (a numeric character from 1 to 4094).
adapter_location	Displays the location of the adapter containing the Ethernet port (any number from 0 to 6).
adapter_port_id	Displays the location of the Ethernet port that is in the adapter (any number from 1 to 4).

Table 41. *Isportip output (continued)*

Attribute	Description
dcbx_state	<p>Displays the DCBx state of the port. A value of:</p> <ul style="list-style-type: none"> • unsupported indicates that the port does not accept Priority Flow Control (PFC) configuration from the switch port, even if the switch is DCBx-capable. All ports that are 1Gbps have this value. • enabled indicates that the connected switch port is enabled for DCBx and the port state is online. • disabled indicates that the connected switch port has disabled for DCBx or the port state is offline. <p>On Ethernet ports that are 10Gb/s, DCBx is automatically enabled if the connected switch port has it enabled.</p> <p>Remember: When this field is disabled or unsupported all fields other than <code>lossless_iscsi</code> and <code>lossless_iscsi6</code> are blank.</p>
iscsi_priority_tag	Displays the numeric priority tag value for the Internet Small Computer System Interface (iSCSI) protocol assigned on the connected switch port. This must be a number from 0 to 7 or blank,
fcoe_priority_tag	Displays the numeric priority tag value for the Fiber Channel over Ethernet (FCoE) protocol assigned on the connected switch port. This value must be a number from 0 to 7 or blank.
pfc_enabled_tags	<p>Displays a list of priority tags for which PFC is enabled on the connected switch port. If you want to use the lossless iSCSI or FCoE function, PFC must be enabled for the corresponding tags on the switch. Once enabled on the switch, the tags are displayed in this field. This value is either blank or a colon-separated list of numbers from 0 to 7.</p> <p>Important: If no priority tags are defined on the switch, this field is blank. If priority tags are defined on the switch but PFC is not enabled for those priority tags, this field is blank.</p>
priority_group_0	Displays the set of priority tags that are within the priority group zero (0). This value is either blank or a colon-separated list of numbers from 0 to 7. This field is part of the Enhanced Transmission Selection (ETS) settings.
priority_group_1	Displays the set of priority tags that are within the priority group one (1). This value is either blank or a colon-separated list of numbers from 0 to 7. This field is part of the ETS settings.
priority_group_2	Displays the set of priority tags that are within the priority group two (2). This value is either blank or a colon-separated list of numbers from 0 to 7. This field is part of the ETS settings.
priority_group_3	Displays the set of priority tags that are within the priority group three (3). This value is either blank or a colon-separated list of numbers from 0 to 7. This field is part of the ETS settings.
priority_group_4	Displays the set of priority tags that are within the priority group four (4). This value is either blank or a colon-separated list of numbers from 0 to 7. This field is part of the ETS settings.
priority_group_5	Displays the set of priority tags that are within the priority group five (5). This value is either blank or a colon-separated list of numbers from 0 to 7. This field is part of the ETS settings.
priority_group_6	Displays the set of priority tags that are within the priority group six (6). This value is either blank or a colon-separated list of numbers from 0 to 7. This field is part of the ETS settings.
priority_group_7	Displays the set of priority tags that are within the priority group seven (7). This value is either blank or a colon-separated list of numbers from 0 to 7. This field is part of the ETS settings.

Table 41. *Isportip output (continued)*

Attribute	Description
bandwidth_allocation	Displays a list of guaranteed bandwidth allocation percentages for priority groups zero (0) through seven (7). The value is either blank or a colon-separated numeric list of bandwidth allocation percentages for each priority group, where each percentage is a whole number integer. This field is part of the ETS settings. Note: The field is blank if no specific bandwidths have been allocated to any priority group on the network.
lossless_iscsi	Displays whether PFC is enabled (on) or not enabled (off) for an iSCSI Internet Protocol Version 4 (IPv4) IP address. To be on: <ul style="list-style-type: none"> It must be a 10Gbps port with a valid IPv4 address PFC for iSCSI is enabled on the switch port Virtual local area network (VLAN) is configured for this IPv4 address iSCSI host attach is enabled on the port Otherwise the value is off.
lossless_iscsi6	Displays whether PFC is enabled (on) or not enabled (off) for an iSCSI Internet Protocol Version 6 (IPv6) IP address. To be on: <ul style="list-style-type: none"> It must be a 10Gbps port with a valid IPv6 address PFC for iSCSI is enabled on the switch port Virtual local area network (VLAN) is configured for this IPv4 address iSCSI host attach is enabled on the port Otherwise the value is off.
storage	Indicates whether the IPv4 address used for iSCSI backend storage attach functions. The values are yes and no (default). If no address is specified the value is blank.
storage_6	Indicates whether the IPv6 address on the port is used for iSCSI backend storage attach functions. The values are yes and no (default). If no address is specified the value is blank.

In the examples below (which list different port configuration options) there are two lines for each possible Ethernet port, which represent the port and iSCSI behavioral effects. Port indices are assigned statically, and higher indices are used for optional ports.

A concise invocation example

```
Isportip -delim ,
```

The resulting output:

```
id,node_id,node_name,IP_address,mask,gateway,IP_address_6,prefix_6,gateway_6,MAC,duplex,state,speed,failover,
link_state,host,remote_copy,host_6,remote_copy_6,remote_copy_status,remote_copy_status_6,vlan,vlan_6,
adapter_location,adapter_port_id,lossless_iscsi,lossless_iscsi6,storage,storage_6
1,1,node1,192.168.48.135,255.255.255.0,192.168.48.1,,,5c:f3:fc:f5:67:ca,Full,configured,1Gb/s,no,
active,yes,1,0,unused,,65,,0,1,off,no,
1,1,node1,,,,,5c:f3:fc:f5:67:ca,Full,configured,1Gb/s,yes,active,,0,,0,,,0,1,,,
2,1,node1,192.168.48.136,255.255.255.0,192.168.48.1,,,5c:f3:fc:f5:67:cb,Full,configured,1Gb/s,no,
active,yes,1,0,unused,,,0,2,off,no,
2,1,node1,,,,,5c:f3:fc:f5:67:cb,Full,configured,1Gb/s,yes,active,,0,,0,,,0,2,,,
3,1,node1,192.168.48.137,255.255.255.0,192.168.48.1,,,00:90:fa:27:ec:22,,configured,10Gb/s,no,
active,yes,1,0,unused,,,1,1,off,no,
3,1,node1,,,,,00:90:fa:27:ec:22,,configured,10Gb/s,yes,active,,0,,0,,,1,1,,,
4,1,node1,192.168.48.138,255.255.255.0,192.168.48.1,0009:2009:0003:0004:0005:0006:0007:1130,64,
fe80:0000:0000:0000:0007:b4ff:fe00:0a00,00:90:fa:27:ec:24,,configured,10Gb/s,no,
active,yes,1,yes,0,unused,,165,170,1,2,on,on,yes,yes
4,1,node1,,,,,00:90:fa:27:ec:24,,configured,10Gb/s,yes,active,,0,,0,,,1,2,,,
1,2,node2,192.168.48.145,255.255.255.0,192.168.48.1,,,5c:f3:fc:f5:68:b2,Full,configured,1Gb/s,no,
active,yes,1,0,unused,,65,,0,1,off,no,
1,2,node2,,,,,5c:f3:fc:f5:68:b2,Full,configured,1Gb/s,yes,active,,0,,0,,,0,1,,,
2,2,node2,192.168.48.146,255.255.255.0,192.168.48.1,,,5c:f3:fc:f5:68:b3,Full,configured,1Gb/s,no,
active,yes,1,0,unused,,,0,2,off,,
2,2,node2,,,,,5c:f3:fc:f5:68:b3,Full,configured,1Gb/s,yes,active,,0,,0,,,0,2,,,
3,2,node2,192.168.48.147,255.255.255.0,192.168.48.1,,,00:90:fa:27:ec:4a,,configured,10Gb/s,no,
active,yes,1,0,unused,,,1,1,off,no,
3,2,node2,,,,,00:90:fa:27:ec:4a,,configured,10Gb/s,yes,inactive,,0,,0,,,1,1,,
```



```
4,2,node2,192.168.48.148,255.255.255.0,192.168.48.1,0009:2009:0003:0004:0005:0006:0007:1230,64,
fe80:0000:0000:0000:0007:b4ff:fe00:0a00,00:90:fa:27:ec:4c,,configured,10Gb/s,no,active,yes,1,yes,
0,unused,,165,170,1,2,on,on,yes,yes
4,2,node2,,,,,,00:90:fa:27:ec:4c,,configured,,yes,inactive,,0,,0,,,,,1,2,,,,,
```

A concise invocation example

```
lsportip
```

The resulting output:

id	node_id	node_name	IP_address	mask	gateway	IP_address_6	prefix_6	gateway_6	MAC	duplex	state
1	1	node1	192.168.1.52	255.255.255.0	192.168.1.1				5c:f3:fc:0b:da:64	Full	config
1	1	node1				fc00:0000:0000:0000:445a:0a17:fcf7:0236	64	fc00:0000:0000:0000:445a:0a17:fcf7:0001	5c:f3:fc:0b:da:66	Full	config
2	1	node1							5c:f3:fc:0b:da:66	Full	config
2	1	node1							5c:f3:fc:0b:da:66	Full	config
1	2	node2	192.168.1.53	255.255.255.0	192.168.1.1				e4:1f:13:2f:b4:a4	Full	config
1	2	node2							e4:1f:13:2f:b4:a4	Full	config
2	2	node2				fc00:0000:0000:0000:445a:0a17:fcf7:0237	64	fc00:0000:0000:0000:445a:0a17:fcf7:0001	e4:1f:13:2f:b4:a6	Full	config
2	2	node2							e4:1f:13:2f:b4:a6	Full	config

A detailed invocation example

```
lsportip 1
```

The detailed resulting output:

```
id 1
node_id 1
node_name node1
IP_address 192.168.20.10
mask 255.255.255.0
gateway 192.168.20.1
IP_address_6
prefix_6
gateway_6
MAC 00:1a:64:97:1b:a0
duplex Full
state online
speed 1Gb/s
failover no
mtu 1500
host yes
remote_copy 0
host_6
remote_copy_6 0
remote_copy_status
remote_copy_status_6
vlan 1063
vlan_6
adapter_location 1
adapter_port_id 1
dcbx_state Enabled
iscsi_priority_tag 4
fcoe_priority_tag 3
pfc_enabled_tags 3:4
pfc_disabled_tags 0:1:2:5:6:7
priority_group_0
priority_group_1
priority_group_2
priority_group_3
priority_group_4
priority_group_5
priority_group_6 3
priority_group_7 4
bandwidth_allocation 0:0:0:0:0:0:30:30
lossless_iscsi on
lossless_iscsi6
storage yes
storage_6

id 1
```

```

node_id 1
node_name node1
IP_address
mask
gateway
IP_address_6
prefix_6
gateway_6
MAC 00:1a:64:97:1b:a0
duplex Full
state online
speed 1Gb/s
failover yes
mtu 1500
host yes
remote_copy 0
host_6
remote_copy_6 0
remote_copy_status
remote_copy_status_6
vlan 1063
vlan_6
adapter_location 1
adapter_port_id 1
dcbx_state Enabled
iscsi_priority_tag 4
fcoe_priority_tag 3
pfc_enabled_tags 3:4
pfc_disabled_tags 0:1:2:5:6:7
priority_group_0
priority_group_1
priority_group_2
priority_group_3
priority_group_4
priority_group_5
priority_group_6 3
priority_group_7 4
bandwidth_allocation 0:0:0:0:0:0:30:30
lossless_iscsi on
lossless_iscsi6
storage
storage_6

```

```

id 1
node_id 2
node_name node2
IP_address 192.168.20.11
mask 255.255.255.0
gateway 192.168.20.1
IP_address_6
prefix_6
gateway_6
MAC 00:1a:64:97:16:08
duplex Full
state online
speed 1Gb/s
failover no
mtu 1500
host yes
remote_copy 0
host_6
remote_copy_6 0
remote_copy_status
remote_copy_status_6
vlan 1063
vlan_6
adapter_location 1

```

```

adapter_port_id 1
dcbx_state Enabled
iscsi_priority_tag 4
fcoe_priority_tag 3
pfc_enabled_tags 3:4
pfc_disabled_tags 0:1:2:5:6:7
priority_group_0
priority_group_1
priority_group_2
priority_group_3
priority_group_4
priority_group_5
priority_group_6 3
priority_group_7 4
bandwidth_allocation 0:0:0:0:0:0:30:30
lossless_iscsi on
lossless_iscsi6
storage yes
storage_6

```

```

id 1
node_id 2
node_name node2
IP_address
mask
gateway
IP_address_6
prefix_6
gateway_6
MAC 00:1a:64:97:16:08
duplex Full
state online
speed 1Gb/s
failover yes
mtu 1500
host yes
remote_copy 0
host_6
remote_copy_6 0
remote_copy_status
remote_copy_status_6
vlan 1063
vlan_6
adapter_location 1
adapter_port_id 1
dcbx_state Enabled
iscsi_priority_tag 4
fcoe_priority_tag 3
pfc_enabled_tags 3:4
pfc_disabled_tags 0:1:2:5:6:7
priority_group_0
priority_group_1
priority_group_2
priority_group_3
priority_group_4
priority_group_5
priority_group_6 3
priority_group_7 4
bandwidth_allocation 0:0:0:0:0:0:30:30
lossless_iscsi on
lossless_iscsi6
storage
storage_6

```

lsportfc

Use the **lsportfc** command to view the status and properties of the Fibre Channel (FC) input/output (I/O) ports for the clustered system.

Syntax

```
lsportfc [-filtervalue --attribute=value] [-filtervalue?] [-nohdr]
          [-delim --delimiter] [object_id]
```

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.

-filtervalue?

(Optional) Displays the valid filter attributes. The following filter attributes for the **lsportfc** command are valid:

- type
- status
- node_id
- fc_io_port_id
- attachment

-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

(Optional) Specifies the ID of an object displayed in the view.

Description

This command enables you to view information about clustered system port status.

Table 42 on page 251 provides the attribute values that can be displayed as output view data.

The following table shows the possible outputs:

Table 42. *Isportfc* output

Attribute	Description
id	Indicates a unique value for the object. The value must be a numeric 0 or greater.
fc_io_port_id	Indicates the FC I/O port ID. The value must be a positive integer.
port_id	Indicates the platform port ID. The value must be a positive integer.
type	Indicates the type of platform port. The value can be either fc or ethernet.
port_speed	Indicates the I/O port speed. The value is XGb. The value is N/A if the port has never been active. If the port is inactive it shows the last-known port speed.
node_id	Indicates the ID of the node containing the port. The value must be a positive integer.
node_name	Indicates the name of the node containing the port.
WWPN	Indicates the I/O port worldwide port name (WWPN). The value must be in 16-character hexadecimal format.
nportid	Indicates the most recent NPort ID used by the port. The value must be in 6-character hexadecimal format, and all zeroes if never active.
status	Indicates if the port is configured to a device of Fibre Channel (FC) port. The values are: <ul style="list-style-type: none"> • active • inactive_configured • inactive_unconfigured.
switch_WWPN	Indicates the WWPN of the device most recently attached to the port. The value must be in 16-character hexadecimal format, or all zeroes if the port has never been active.
fpma	Indicates the Fabric Provided MAC Address (FPMA) assigned to the Fiber Channel over Ethernet (FCoE) VN_PORT port. The value is a formatted 48-bit MAC address. If the switch is connected to an FCF, the value is N/A for ports that are never active. If the port is inactive the last-known FPMA is used.
vlan_id	Indicates the VLAN ID on which a specific VN port is communicating. The value is up to a 4-character decimal string. The value is N/A for ports that are never active. If the port is inactive the last-known VLAN ID is used.
fcf_MAC	Indicates the MAC address for the switch attached to the VN port. The value is N/A for ports that are never active. The value is a formatted 48-bit MAC address. If the port is inactive the last known fcf_MAC value is used.
attachment	Indicates if the port is attached to an FC switch or directly to an FC host. (Or, if the port is offline, it specifies which one it was attached to when last online.)
cluster_use	Indicates the node's current capability for local or partner cluster communications: <ul style="list-style-type: none"> • none indicates the port cannot be used for any node communication • local indicates the port can be used for local clustered system (system) node communication • partner indicates the port can be used for partner system node communication • local_partner indicates the port can be used for both local and partner system node communication
adapter_location	Indicates the location of the adapter containing the Ethernet port (any number from 0 to 6).
adapter_port_id	Indicates the location of the Ethernet port that is in the adapter (any number from 1 to 4).
fabric_WWN	Indicates the World Wide Name (WWN) value for the connected fabric string. The value must be in 16-character hexadecimal format, or all zeroes if the port has never been active or the port is not connected to a fabric.

A concise invocation example

```
lsportfc
```

The resulting output:

id	fc_io_port_id	port_id	type	port_speed	node_id	node_name	WWPN	nportid	status	attachment	cluster_use	adapter_location	adapter_port_id
0	1	1	fc	8Gb	1	node1	500507680140BADD	0E2411	active	switch	local_partner	1	1
1	2	2	fc	8Gb	1	node1	500507680130BADD	0E2412	active	switch	local_partner	1	2
2	3	3	fc	N/A	1	node1	500507680110BADD	000000	inactive_unconfigured	none	partner	1	3
3	4	4	fc	N/A	1	node1	500507680120BADD	000000	inactive_unconfigured	none	none	1	4
4	5	3	ethernet	10Gb	1	node1	500507680150BADD	0E2413	active	switch	local	2	1
5	6	4	ethernet	10Gb	1	node1	500507680160BADD	0E2414	inactive_configured	switch	local	2	2
6	1	1	fc	N/A	2	node2	500507680140BADE	000000	inactive_unconfigured	none	local_partner	2	3
7	2	2	fc	N/A	2	node2	500507680130BADE	000000	inactive_unconfigured	none	local_partner	2	4
8	3	3	fc	N/A	2	node2	500507680110BADE	000000	inactive_unconfigured	none	partner	3	1
9	4	4	fc	N/A	2	node2	500507680120BADE	0E2414	active	switch	none	3	2
10	5	3	ethernet	10Gb	2	node2	500507680150BADE	0E2415	active	switch	local	3	3
11	6	4	ethernet	10Gb	2	node2	500507680160BADE	0E2416	active	switch	local	3	4

A detailed invocation example

```
lsportfc 10
```

The detailed resulting output:

```
id 10
fc_io_port_id 5
port_id 3
type ethernet
port_speed 10Gb
node_id 6
node_name node3
WWPN 50050768015051E5
nportid 012701
status active
switch WWPN 202700053346FA3D
fpma 0E:FC:00:01:27:01
vlanid 100
fcf_MAC 00:05:73:C2:CA:B4
cluster_use none
adapter_location 1
adapter_port_id 1
fabric_WWN 202700053346FA3C
```

Isportsas

Use the **lsportsas** command to display the status of all SAS ports in the clustered system.

Syntax

```
➤➤ lsportsas — [ -filtervalue —attribute=value ] [ -filtervalue? ] [ -nohdr ] ➤➤
➤ [ -delim —delimiter ] ➤➤
```

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.

-filtervalue?

(Optional) Displays the valid filter attributes. The following filter attributes for the **lsportsas** command are valid:

- node_id

- status
- attachment
- type

-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 view. This 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 allowable width for each data item. 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 enables you to view information about system port status.

This command output shows all available paths, defined by zoning, independent of their usage. This means the command output includes paths not used because of port masking.

Table 43 provides the attribute values that can be displayed as output view data.

Table 43. lsportsas output

Attribute	Description
id	Indicates the line number within the displayed information (numeric string).
port_id	Indicates the ID of the port.
port_speed	Indicates the speed of the I/O port (in XGb). This is the fastest local link speed for the SAS port. The value is the last-known port speed if the port is inactive, and N/A if port has is unused and has never been active.
node_id	Indicates the ID of the node that contains the port (numeric string).
node_name	Indicates the name of the node that contains the port (alphanumeric string).
WWPN	Indicates the worldwide port name (WWPN) for the I/O port (16-character hexadecimal string).
status	Indicates the status of the port (numeric string). The values can be: <ul style="list-style-type: none"> • online if functioning • offline if not functioning • offline_unconfigured if not functioning but not configured by the user • degraded if one or more ports are not functioning or have a lower speed than others • excluded if excluded by the user or system
switch_WWPN	Indicates the WWPN of the switch port if attached to switch (16-character hexadecimal string), or is blank. If the port is offline, the last-known value is displayed.

Table 43. *lsportsas* output (continued)

Attribute	Description
attachment	Indicates what the port is attached to. The possible values are: <ul style="list-style-type: none"> • drive • controller • switch • host • enclosure • none If the port is offline, the field shows what was attached when the port was last online.
type	Indicates how the port is configured. This field also shows the devices that can be attached to the SAS port. The possible values are: <ul style="list-style-type: none"> • drive • enclosure • enclosure_controller • host_controller • none
adapter_location	Displays the location of the adapter containing the Ethernet port (any number from 0 to 6).
adapter_port_id	Displays the location of the Ethernet port that is in the adapter (any number from 1 to 4).

An invocation example

`lsportsas`

The resulting output:

id	port_id	port_speed	node_id	node_name	WWPN	status	switch_WWPN	attachment	type	adapter_location	adapter_port_id
0	1	3Gb	1	node1	500507680140004A	offline		enclosure	enclosure	0	1
1	2	6Gb	1	node1	500507680150004A	online	5001234567892000	switch	host_controller	0	2
4	1	3Gb	2	node2	50050768014051E5	online		host	host_controller	0	3
5	2	3Gb	2	node2	50050768015051E5	offline_unconfigured		none	none	0	4

An invocation example

`lsportsas`

The resulting output:

id	port_id	port_speed	node_id	node_name	WWPN	status	switch_WWPN	attachment	type	adapter_location	adapter_port_id
0	1	12Gb	1	node1	50050768056C009E	online	500507680600B63F	enclosure	enclosure	0	0
1	2	12Gb	1	node1	50050768056C009F	online	500507680600B64F	enclosure	enclosure	0	1
2	0	12Gb	1	node1	50050768056C009G	online	500507680600B65F	enclosure	internal	0	2
4	1	12Gb	2	node2	50050768056C009I	online	500507680600B66F	enclosure	enclosure	0	3
5	2	12Gb	2	node2	50050768056C009J	online	500507680600B67F	enclosure	enclosure	0	4
6	0	12Gb	2	node2	50050768056C009K	online	500507680600B68F	enclosure	internal	1	1

lsquorum

Use the **lsquorum** command to list the quorum devices that the clustered system (system) is currently using to store quorum data.

Syntax

```

>>> 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 device by its index number. When you use this parameter, a detailed view of the specified device is returned. If you do not specify a device, then a concise view of all quorum devices is displayed.

Description

This command displays a concise list or a detailed view of the MDisk or drives that the system 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, but only MDisks are used to hold quorum data. If the quorum object type is a drive, the controller ID and name fields are blank.

Table 44 provides the attribute values that can be displayed as output view data.

Table 44. lsquorum output

Attribute	Possible Values
quorum_index	Indicates the quorum device by index number.
status	Indicates the quorum device status.
name	Indicates the name of the object used as the quorum device.
controller_id	Indicates the ID of the controller of an mdisk object used as the quorum device.
controller_name	Indicates the name of the controller of an mdisk object used as the quorum device.
active	Indicates if this is the active quorum device the system will use as a tie breaker.
object_type	Indicates the type of object the quorum device uses.
override	Indicates if the automatic quorum selection for this quorum device was overridden.
site_id	Indicates the site value for the quorum device. This numeric value is 1, 2, 3 or blank.

Table 44. *lsquorum* output (continued)

Attribute	Possible Values
site_name	<p>Indicates the site name for the quorum device (MDisks or drives). This is an alphanumeric value or is blank.</p> <p>For the quorum application this identifies the site it is deployed at. By default, this is the local host's IP address, but a custom alias can also be set.</p>

A concise invocation example

```
lsquorum
```

The concise resulting output:

quorum_index	status	id	name	controller_id	controller_name	active	object_type	override	site_id	site_name
0	online	1	mdisk1	1	controller1	no	mdisk	no	2	site2
1	online	2	mdisk2	1	controller1	no	mdisk	no	1	site1
2	online					yes	device	no		quorumhost/9.15

A detailed invocation example

```
lsquorum 1
```

The detailed resulting output:

```
quorum_index 1
status online
id 309
name mdisk9
controller_id 1
controller_name controller3
active yes
object_type drive
override yes
site_id 1
site_name CPD1

quorum_index 2
status online
id 33
name
controller_id
controller_name
active no
object_type drive
override no
site_id 1
site_name CPD1
```

lsroute

Use the **lsroute** command to display the IP routing table.

Syntax

```

➤➤ lsroute — [ -delim — delimiter ] [ -nohdr ]

```

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 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 **lsroute** command is equivalent to the Linux **route** command.

An invocation example

```
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

lstimezones

Use the **lstimezones** command to list the time zones that are available on the clustered system (system). Each timezone is assigned an ID that can be used in the **settimezone** command to set the time zone.

Syntax

```
➤➤ — 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 system. Each time zone is assigned an ID. This ID can be used in the **settimezone** command.

An invocation example

```
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
```

Issasportcandidate

Use the **lssasportcandidate** command to list the unconfigured serial-attached SCSI (SAS) ports logged in and available to add to the SAS worldwide port name (WWPN) or host objects.

Syntax

```
➤ lssasportcandidate — [ -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 SAS ports.

Note: The **lssasportcandidate** command presents a list of host SAS ports that are logged in to nodes. However, there are situations when the list might include host SAS ports that are no longer logged in or even part of the SAN fabric. For example, if a host SAS port is unplugged from a switch but **lssasportcandidate** shows the WWPN that is logged in to all nodes, the incorrect entry is removed when another device is plugged in to the same switch port that previously contained the removed host SAS port.

This shows the possible output:

Table 45. *lssasportcandidate* output

Attribute	Description
sas_WWPN	Indicates that the SAS WWPN that is logged in is unconfigured (not assigned to a host). This value must be 16 hexadecimal characters.

An invocation example

```
lssasportcandidate
```

The resulting output:

```
sas_WWPN
200600A0B813B7AC
200600A0B813B7AD
```

lsssecurity

Use the **lsssecurity** command to display the current system Secure Sockets Layer (SSL) or Transport Layer Security (TLS) security settings.

Syntax

```
►► lsssecurity [ -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 current system SSL or TLS security settings.

This table provides the possible values displayed for the **lssecurity** command.

Table 46. *lssecurity* attribute values

Attribute	Value
sslprotocol	<p>Identifies the current security level setting, a numeric value of 1, 2, or 3.</p> <p>A security level setting of:</p> <ul style="list-style-type: none"> • 1 allows TLS 1.0, TLS 1.1, and TLS 1.2, but disallows SSL 3.0 • 2 disallows TLS 1.0 and TLS 1.1 • 3 additionally disallows TLS 1.2 cipher suites that are not exclusive to 1.2 <p>Note: You cannot use the management GUI if the sslprotocol value is set to 1 and you are using SSL 3.0 or TLS 1.0.</p>

An invocation example

```
lssecurity
```

The resulting output:

```
sslprotocol:1
```

Issite

Use the **lsite** command to report the names of the sites.

Syntax

```

>> lsite [ -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 reports the names of the sites.

Remember: This command is only applicable when a system is configured as a stretched system or a HyperSwap system (using the **chsystem -topology** command).

In a stretched configuration these applications are spread across two or more geographic locations or sites:

- Nodes
- Storage
- Host servers
- Infrastructure

Table 47 provides the attribute values that can be displayed as output view data.

Table 47. *lssite attribute values*

Attribute	Value
id	Identifies the numeric value representing the site. The value can be 1, 2, or 3.
name	Identifies the site name.

An invocation example

```
lssite
```

The resulting output:

```
id name
1  CPD1
2  CPD2
3  Quorum
```

lsthrottle

Use the **lsthrottle** command to list throttle objects configured in the clustered system.

Syntax

```
➤ lsthrottle - [ -nohdr ] [ -filtervalue? ] [ -filtervalue attribute=value ]
➤ [ -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 no data is displayed, headings are not displayed.

-filtervalue?

Displays a list of valid filter attributes for the **-filtervalueattribute=value** parameter. The valid filters for the **lsthrottle** command are:

- throttle_type

-filter *value 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 when you are using the CLI:

- The wildcard character is an asterisk (*).
- The command can contain a maximum of one wildcard.
- When you are using a wildcard, you must enclose the filter entry within double quotation marks ("").

-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 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 1-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 lists throttle objects configured in the clustered system.

Table 48 provides the attribute values that can be displayed as output view data.

Table 48. lsthrottle output

Attribute	Description
throttle_id	Indicates the unique ID for the throttle object. The value is a number 0 - 10144.
throttle_name	Indicates the unique name for the throttle object. The value is an alphanumeric string that is 63 characters long.
object_id	Indicates the ID of the object to which throttle is applied. The value is a number 0 - 8191.
object_name	Indicates the name of the object to which throttle is applied. The value is an alphanumeric string that is 63 characters long.
throttle_type	Indicates the type of throttle object. The values are: offload, vdisk, host, hostcluster, and mdiskgrp.
IOPs_limit	Indicates the limit of configured IOPs. The value is a numeric string 0 - 33554432. If no limit is specified the value is blank.
bandwidth_limit_MB	Indicates the bandwidth (in MBps). The value is a numeric string 0 - 268435456. If no limit is specified the value is blank.

An invocation example

`lsthrottle`

The detailed resulting output:

throttle_id	throttle_name	object_id	object_name	throttle_type	IOPs_limit	bandwidth_limit_MB
0	throttle0	1	R48U20_213	host		40
1	throttle1	0	WinHostClust	hostcluster	8000	
2	throttle2	9	vdisk0	vdisk		20
3	throttle3	11	mdiskgrp0	mdiskgrp		100
0	throttle4			offload		500

Issystem

Use the **Issystem** command to display a detailed view of a clustered system (system).

Syntax

```
►► Issystem — [ -nohdr ] [ -delim — delimiter ] ►►
```

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) 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 character with one byte. In a detailed view, the data is separated from its header by the specified delimiter.

Description

This command displays a detailed view of a system.

Table 49 provides the attribute values that can be displayed as output view data.

Table 49. Issystem output

Attribute	Possible Values
layer	The value can be either: <ul style="list-style-type: none">• replication, which means the system can create partnerships• storage (default), which means the system can present storage
location	The location is <code>local</code> or <code>remote</code> .
statistics status	The status is <code>on</code> or <code>off</code> .
auth_service_type	Native Lightweight Directory Access Protocol (LDAP)
auth_service_configured	True if the <code>auth_service_type</code> is configured and either these is true: <ul style="list-style-type: none">• The <code>auth_service_type</code> is LDAP-only (if at least one LDAP server is configured)• The <code>auth_service_type</code> is only IBM Security Services (referred to as TIP in the CLI):<ul style="list-style-type: none">– The name, password, and URL are established– An SSL certificate is created (if an HTTPS URL is available)
auth_service_enabled	True if the <code>auth_service_type</code> is configured.

Table 49. *Issystem output (continued)*

Attribute	Possible Values
email_state	The possible values are: <ul style="list-style-type: none"> • running • stopped • invalid
partnership	The possible values are: <ul style="list-style-type: none"> • fully_configured • 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 • Blank
tier	Indicates which system information is being reported. The values are: <ul style="list-style-type: none"> • tier0_flash • tier1_flash • tier_enterprise • tier_nearline
tier_capacity	Indicates the total MDisk storage in the tier.
tier_free_capacity	Indicates the amount of MDisk storage in the tier that is unused.
compression_active	Indicates whether there are any compressed volume copies in the system.
compression_virtual_capacity	Indicates the total virtual capacity for all compressed volume copies in the system. This value is in unsigned decimal format.
compression_compressed_capacity	Indicates the total used capacity for all compressed volume copies in the system. This value is in unsigned decimal format.
compression_uncompressed_capacity	Indicates the total uncompressed used capacity for all compressed volume copies in the system. This value is in unsigned decimal format.
rc_buffer_size	Indicates the resource buffer size that is assigned for Metro Mirror, Global Mirror, and HyperSwap Copy Services.
has_nas_key	The value is yes or no.
total_drive_raw_capacity	The total known capacity of all discovered drives (regardless of drive use)
email_organization	Indicates the user's organization that is shown in the Call Home email function.
email_machine_address	Indicates the user's mailing address that is shown in the Call Home email function.

Table 49. *Issystem output (continued)*

Attribute	Possible Values
email_machine_city	Indicates the user's city that is shown in the Call Home email function.
email_machine_state	Indicates the user's state that is shown in the Call Home email function.
email_machine_zip	Indicates the user's zip code or postal code that is shown in the Call Home email function.
email_machine_country	Indicates the user's country that is shown in the Call Home email function.
cache_prefetch	Indicates whether cache prefetching is enabled across the system. The values are on and off.
local_fc_port_mask	Indicates the Fibre Channel (FC) input/output (I/O) ports that a system can use for node-to-node communications on a local system if those FC I/O ports exist on a node. The value is 64 binary bits.
partner_fc_port_mask	Indicates the FC I/O ports that a system can use for system-to-system communications on a partner system if those FC I/O ports exist on a node. The value is 64 binary bits.
topology	Indicates the system topology: <ul style="list-style-type: none"> • standard • stretched • hyperswap
topology_status	Indicates the system topology status: <ul style="list-style-type: none"> • dual_site • recovered_site_1 • recovered_site_2
rc_auth_status	Indicates remote copy authentication. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap. The values are: <ul style="list-style-type: none"> • chap • none (default)
vdisk_protection_time	Indicates the volume protection time in minutes (whether volume protection is enabled or not enabled). The value must be a number from 15 (default) to 1440.
vdisk_protection_enabled	Indicates whether volume protection is enabled (yes) or disabled (no) for a system.
product_name	Indicates the product name (an alphanumeric string of no more than 62 characters).
odx	Indicates whether offloaded data transfer (ODX) is enabled or disabled. The values are on and off.
easy_tier_acceleration	Indicates Easy Tier and pool balancing acceleration status. The values are on and off.
max_replication_delay	Indicates the value for maximum replication delay. This is a numeric value from 0 to 360.
partnership_exclusion_threshold	Indicates the partnership exclusion threshold value. This is a numeric value from 30 to 315.

Table 49. *Issystem output (continued)*

Attribute	Possible Values
gen1_compatibility_mode	Indicates whether a system is compatible with a downlevel system (such as a Storwize V7000 Gen1) system. The values are yes and no.
ibmcustomer	Indicates the customer number. The value is blank or is a number that contains 7 - 10 digits.
ibmcomponent	Indicates the component. The value is blank or SANVCNSW1.
ibmcountry	Indicates the country. The value is blank or a 3-digit number.
tier_0_flash_compressed_data_used	Indicates the capacity of compressed data used on the flash tier 0 storage tier (in MB). The value must be a number with two decimal places.
tier_1_flash_compressed_data_used	Indicates the capacity of compressed data used on the flash tier 1 storage tier (in MB). The value must be a number with two decimal places.
tier_enterprise_compressed_data_used	Indicates the capacity of compressed data used on the tier 2 enterprise storage tier (in MB). The value must be a number with two decimal places.
tier_nearline_compressed_data_used	Indicates the capacity of compressed data used on the tier 3 nearline storage tier (in MB). The value must be a number with two decimal places.

Information about the remote system is reported by the **Issystem** command if you issue either the **mkfcpartnership** or **mkippartnership** command from the local system to the remote system. For example, if the partnership is at least partially established from the local system.

Issue the **Issystem** command to display a detailed view of the system.

Detailed view shows the fields that are described for remote systems only; if the system **location** is **local**, then **partnership** and **bandwidth** do not apply (and are not defined or provided). For a remote system, these fields indicate the following information:

- **location:** remote
- **partnership:**

fully_configured

The **mkfcpartnership** or **mkippartnership** command is issued in both directions and the remote system is online and available.

partially_configured_local

The **mkfcpartnership** or **mkippartnership** command is issued from the local system to the remote system. The remote system is online and available for partnership.

partially_configured_local_stopped

The **mkfcpartnership** or **mkippartnership** command is issued from the local system to the remote system. The **chpartnership** command with the **stop** parameter is issued from the local system, and the remote system is online and available. Issue the **chpartnership** command with the **start** parameter on the local system, and **mkfcpartnership** or **mkippartnership** on the remote system.

not_present

The **mkfcpartnership** or **mkippartnership** command is issued from the local system to the remote system, and the remote system is not available. Either the remote system is offline, or it is not connected to the local system.

fully_configured_stopped

The **mkfcpartnership** or **mkippartnership** command is issued in both directions and the remote system is online and available. The **chpartnership** command with the **stop** parameter is issued from the local system.

fully_configured_remote_stopped

The **mkfcpartnership** or **mkippartnership** command is issued in both directions and the remote system is online and available. The **chpartnership** command with the **stop** parameter is issued from the remote system.

fully_configured_local_excluded

The **mkfcpartnership** or **mkippartnership** command is issued in both directions. The local system is excluding the connection to the remote system due to too many problems, or either system in the partnership is unable to sustain the I/O workload for the Metro Mirror, Global Mirror, or HyperSwap relationships.

fully_configured_remote_excluded

The **mkfcpartnership** or **mkippartnership** command has been issued in both directions. The remote system is excluding the connection to the local system due to too many problems, or either system in the partnership is unable to sustain the I/O workload for the Metro Mirror, Global Mirror, or HyperSwap relationships.

fully_configured_exceeded

There are too many systems in the system network, and the partnership from the local system to the remote is disabled. Refer to the 1710 or 1720 errors in the system error log at the local and remote system.

- **bandwidth:** The bandwidth available on the intersystem link for background copy, in megabytes per second (MBps).

The `console_IP` field displays either the:

- Automatically-populated system port 1 IP Address - Internet Protocol Version 4 (IPv4) or IPv6
- User-populated IPv4 address

The port value is always `443`, which requires the system to run by using default Hypertext Transfer Protocol Secure (HTTPS).

A detailed invocation example

```
lssystem -delim :
```

The resulting output:

```
id:00000200A2600906
name:tbcluster-29
location:local
partnership
bandwidth
total_mdisk_capacity:60.5TB
space_in_mdisk_grps:60.5TB
space_allocated_to_vdisks:643.74GB
total_free_space:59.9TB
total_vdiskcopy_capacity:663.46GB
total_used_capacity:560.99GB
total_overallocation:1
total_vdisk_capacity:501.25GB
total_allocated_extent_capacity:792.50GB
statistics_status:on
statistics_frequency:15
cluster_locale:en_US
time_zone:375 Europe/London
code_level:6.4.0.0 (build 64.6.1205081000)
console_IP:9.71.53.69:443
id_alias:00000200A2600906
```



```
max_replication_delay:222
partnership_exclusion_threshold:120

ibmcustomer:1234567
ibmcomponent:SANVCNSW1
ibmcountry:866
tier_0_flash_compressed_data_used:0.00MB
tier_1_flash_compressed_data_used:0.00MB
tier_enterprise_compressed_data_used:0.00MB
tier_nearline_compressed_data_used:0.00MB
```

A detailed invocation example

lssystem

The resulting output:

```
id 000002006C40A278
name cluster0
location local
partnership
bandwidth
total_mdisk_capacity 222.2GB
space_in_mdisk_grps 0
space_allocated_to_vdisks 0.00MB
total_free_space 222.2GB
total_vdiskcopy_capacity 0.00MB
total_used_capacity 0.00MB
total_overallocation 0
total_vdisk_capacity 0.00MB
total_allocated_extent_capacity 0.00MB
statistics_status on
statistics_frequency 15
cluster_locale en_US
time_zone 522 UTC
code_level 6.4.0.0 (build 61.9.1112130001)
console_IP 0.0.0.0:443
id_alias 000002006C40A278
gm_link_tolerance 300
gm_inter_cluster_delay_simulation 0
gm_intra_cluster_delay_simulation 0
gm_max_host_delay 5
email_reply
email_contact
email_contact_primary
email_contact_alternate
email_contact_location
email_contact2
email_contact2_primary
email_contact2_alternate
email_state stopped
inventory_mail_interval 0
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
auth_service_type tip
relationship_bandwidth_limit 25

tier tier0_flash
tier_capacity 1.63TB
tier_free_capacity 1.63TB
```

```
ibmcustomer 1234567
ibmcomponent SANVCNSW1
ibmcountry 866
tier_0_flash_compressed_data_used:0.00MB
tier_1_flash_compressed_data_used:0.00MB
tier_enterprise_compressed_data_used:0.00MB
tier_nearline_compressed_data_used:0.00MB
```


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, with the width of each column set to the maximum width of each item of data. In a detailed view, each item of data is an individual row, and if you display headers, the data is separated from the header by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a 1-byte character. Enter **-delim :** on the command line, and 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 specified *delimiter* separates the data from its header.

Description

This command lists information about the current system Secure Sockets Layer (SSL) certificate and indicates if there is an outstanding certificate request.

Table 50 provides the attribute values that can be displayed as output view data.

Table 50. Issystemcert output

Attribute	Possible Values
certificate	Displays a readable version of the current SSL certificate.
certificate export	Displays an encoded version of the SSL certificate.
certificate_request_outstanding	Indicates whether there is a certificate request that has not been completed by installing the signed certificate. The value is yes or no.

An invocation example

```
lssystemcert
```

The detailed resulting output:

```
certificate: 58 fields
```

```
Data:
```

```
Version: 3 (0x2)
```

```
Serial Number: 1431938814 (0x5559a6fe)
```

```
Signature Algorithm: sha256WithRSAEncryption
```

```
Issuer: C=US, L=Springfield, O=TMI, OU=ABC, CN=2154/emailAddress=chili@snpp.com
```

```
Validity
```

```
Not Before: May 18 08:46:54 2015 GMT
```

```
Not After : May 14 08:46:54 2030 GMT
```

```
Subject: C=US, L=Springfield, O=TMI, OU=ABC, CN=2154/emailAddress=chili@snpp.com
```

```
Subject Public Key Info:
```

```
Public Key Algorithm: rsaEncryption
```

```
Public-Key: (2048 bit)
```

```
Modulus:
```

```
00:de:1c:70:c2:91:87:3c:6a:92:91:f7:d9:a3:5b:
05:e6:91:f1:87:c1:25:38:61:ad:4d:d9:26:19:7b:
9e:61:a5:fd:b1:d1:eb:d1:e4:a8:78:21:75:58:80:
4a:5c:dd:5e:6c:8b:1b:de:57:f9:d5:1f:71:92:3e:
78:d5:a4:75:1e:11:b2:62:18:52:0f:4d:32:a8:fd:
2b:16:4f:42:d1:d6:70:af:86:eb:fe:a1:ab:bc:66:
8a:44:bc:e0:36:53:77:96:2f:74:7d:95:33:79:c2:
59:5e:e1:43:50:da:43:25:c4:5d:3a:ac:d7:82:ad:
34:d5:ba:4c:52:4a:c0:81:3a:ad:e8:33:fe:4f:be:
e8:47:fa:5b:1f:dd:d8:9e:3b:44:a6:b6:b9:43:d2:
d4:45:8e:cb:5b:bb:10:5b:c9:30:68:2c:30:b6:e4:
ea:59:6d:a2:37:a7:13:77:28:1d:13:68:58:7b:dd:
90:d6:a8:81:7b:79:9f:1e:e4:a7:67:1b:7b:c5:b4:
90:dc:6b:d4:1f:7e:e9:e3:7b:ac:26:59:11:f1:99:
34:f0:6a:50:41:76:ad:a3:30:74:8f:8f:f5:ed:1e:
```

```

21:77:ff:51:90:1b:83:fb:04:f0:62:3d:71:17:a5:
ab:44:e8:bc:b0:82:0d:af:af:ae:68:5a:cf:e3:c8:
a9:53
Exponent: 65537 (0x10001)
X509v3 extensions:
  X509v3 Basic Constraints:
    CA:FALSE
  Netscape Comment:
    OpenSSL Generated Certificate
  X509v3 Subject Key Identifier:
    87:66:33:16:61:7A:8E:CA:B4:BA:78:7B:56:56:8A:9D:C5:96:80:76
  X509v3 Authority Key Identifier:
    keyid:87:66:33:16:61:7A:8E:CA:B4:BA:78:7B:56:56:8A:9D:C5:96:80:76

Signature Algorithm: sha256WithRSAEncryption
56:b1:5d:59:11:ae:7b:6e:29:cc:1f:a8:75:77:d2:65:d6:88:
75:8e:b9:cd:d6:71:ac:7e:89:8c:65:68:36:a8:28:97:88:36:
42:da:a4:58:9b:c6:ce:c1:56:c9:0e:c5:ce:e7:01:74:d0:66:
d0:4d:d3:0f:84:53:f6:e5:89:8e:44:6d:70:13:45:9c:21:91:
50:f4:b0:b7:cc:cb:18:e8:d7:b3:38:b4:f5:5d:36:51:8c:7e:
52:d4:24:0f:1f:2e:0a:b4:b6:9b:cb:23:43:6c:16:a2:a5:de:
84:8a:0d:28:3c:d9:3d:5d:a4:52:44:28:90:98:a6:26:a9:c9:
87:6c:27:3f:ef:09:5f:9d:0b:40:8d:07:64:ee:33:d9:40:47:
98:02:10:58:2b:54:33:d9:37:69:d4:13:e6:0d:ec:46:26:b1:
c1:c5:15:7c:8d:89:26:f7:95:d9:2f:d9:33:8c:f0:1a:dc:08:
19:eb:18:16:51:30:a3:c0:ee:be:86:7d:3d:91:61:d5:99:bf:
5e:19:b9:89:72:e1:4c:ea:5e:2b:90:ce:ce:75:83:e0:c9:14:
83:21:21:e0:f8:28:94:90:71:e6:13:ca:97:8c:e3:58:b9:0c:
62:03:e5:1c:1b:6c:dd:c3:60:48:d4:78:24:8e:22:34:78:32:
fe:45:ee:36
certificate_export: 23 fields
-----BEGIN CERTIFICATE-----
MIIDzTCCArWgAwIBAgIEVmm/jANBgkqhkiG9w0BAQsFADBqMQswCQYDVQGEwJH
QjEQA4GA1UEBwHSHVyc2xleTEEMMAoGA1UECgwDSUJNMWwCgYDVQQLDANTU0cx
DTALBgNVBAMBDIxNDUxHjAcBgkqhkiG9w0BCQEWDD3N1cHBvcnRAaWJtLmNvbTAE
Fw0xNTA1MTgwODQ2NTRaFw0zMDA1MTQwODQ2NTRaMGoxCzAJBgNVBAYTAkdCMRAw
DgYDVQQHDAdIdXZjbGV5MQwwCgYDVQQKDANJQk0xDDAKBgNVBAsMAINTRzENMAAG
A1UEAwEMjE0NTEeMBwGCSqGSIb3DQEJARYPc3VwcG9ydEBpYm0uY29tMIIjAN
BgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEA3hxwwpGHPGqSkffZo1sF5pHxh8E1
OGGtTdkmGXueYaX9sdHr0eSoeCF1WIBKXN1ebIsb31f51R9xkj541aR1HhGyYhhS
D00yqP0rFk9C0dZwr4br/qGrvGaKRLzgN1N31i90fZUzecJZXuFDUNpDJcRd0qzX
gg001bpMukrAgTqt6DP+T77oR/pbH93YnjtEpra5Q9LURY7LW7sQW8kwaCwwtuTq
WW2iN6cTdygdE2hYe92Q1qiBe3mfHuSnZxt7xbSQ3GvUH37p43usJ1kR8Zk08GpQ
QXatozB0j4/17R4hd/9RkBuD+wTwYj1xF6Wrr0i8sIINr6+uaFrP48ipUwIDAQAB
o3sweTAJBgNVHRMEAjAAMCwGCWGSAGG+EIBDQ0FFh1PcGVuU1NMIEdlbnV5YXR1
ZCBBDZXJ0aWZpY2F0ZTAdBgNVHQ4EFgQUh2YzFmF6jsq0unh7V1aKncWwGHyHwYD
VR0jBBgwFoAUh2YzFmF6jsq0unh7V1aKncWwGHyHwYDQYJKoZIhvcNAQELBQADggEB
AFaxXVkrRntuKcwfqHV30mXWiHW0uc3Wcax+iYx1aDaoKJeINKLapFibxs7BVsk0
xc7nAXTZtBN0w+EU/b1iY5EbXATRZwhkVD0sLfMyxjo17M4tPvDN1GMf1LUJA8f
Lgq0tpvLI0NsFqK13oSKDSg82T1dpFJEKJCYpiapyYdsJz/vCV+dC0CNB2TuM91A
R5gCEGrVDPZN2nUE+YN7EYmscHFFXyNiSb31dkv2TOM8BrCBnrGBZRMKPA7r6G
fT2RYdWzV14ZuY1y4UzqXiUQzs51g+DJFIMhIeD4KJSQceYTypeM41i5DGID5Rwb
bn3DYEjUeCS0IjR4Mv5F7jY=
-----END CERTIFICATE-----
certificate_request_outstanding:no

```

Issystemip

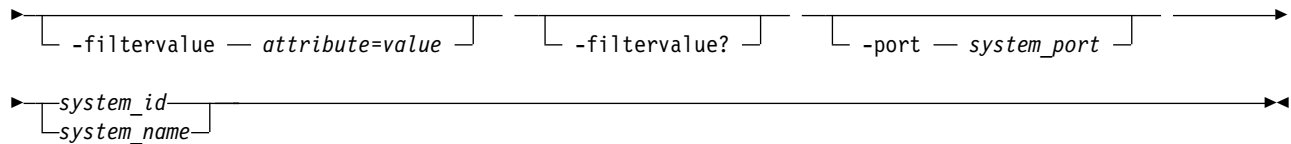
Use the **Issystemip** command to display a list of the clustered system (system) management IP addresses configured for each port.

Syntax

```

►► Issystemip — [ -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 data item 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 when using the command-line interface (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 shown in the following example:

```
lssystemip -filtervalue "system_name=md*"
```

-filtervalue?

(Optional) displays a list of filters that can be applied against this view. The following filter attributes are valid for the **lssystemip** command:

- port_id
- system_name
- system_id

system_id | system_name

(Required) Specifies the name or ID of a system.

-port system_port

(Required) Specifies the system port (1 or 2) to apply changes to.

Description

This command displays a list of the system management IP addresses configured for each port.

Table 51 on page 274 provides the attribute values that can be displayed as output view data.

Table 51. *lssystemip* output

Attribute	Possible Values
cluster_id	Indicates the ID of the system.
cluster_name	Indicates the name of the system.
location	Indicates the system location.
port_id	Indicates the ID of the port.
IP_address	Indicates the Internet Protocol Version 4 (IPv4) address.
subnet_mask	Indicates the IPv4 subnet mask.
gateway	Indicates the IPv4 gateway.
IP_address_6	Indicates the Internet Protocol Version 6 (IPv6) address.
gateway_6	Indicates the IPv6 gateway.
prefix_6	Indicates the IPv6 prefix.

A concise invocation example

```
lssystemip -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

```
lssystemip 000002006CC0B71A
```

The detailed resulting output:

```
cluster_id 000002006CC0B71A
cluster_name c11
location local
port_id 1
IP_address 192.168.1.2
subnet_mask 255.255.255.0
gateway 192.168.1.1
IP_address_6 2001:0db8:85a3:0000:0000:8a2e:0370:7334
gateway_6 2001:0db8:85a3:0000:0000:8a2e:0370:7334
prefix_6 64


cluster_id 000002006CC0B71A
cluster_name c11
location local
port_id 2
IP_address 192.168.1.2
subnet_mask 255.255.255.0
```

```
gateway 192.168.1.1
IP_address_6 2001:0db8:85a3:0000:0000:8a2e:0370:7334
gateway_6 2001:0db8:85a3:0000:0000:8a2e:0370:7334
prefix_6 64
```

lssystemstats

Use the **lssystemstats** command to display the most recent values of all node statistics in a clustered system (system), or to display a history of values for a given subset of available statistics across all nodes in a system. This command also can be used to display a history of values for a given subset of available statistics.

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.

-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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):

```
lssystemstats -filtervalue stat_name="io*"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue attribute=value** parameter:

- **stat_name**

-history stat_list

Provides the most recent node statistical values, specific node statistical values, or historical data for any node.

Description

This command returns one set of statistics for all the nodes in the system. The statistical values are determined using samples received from each node .

Note: Values are rounded to the nearest integer when appropriate (for example, between one and ninety-nine when considering percentages).

Table 52 provides the attribute values that can be displayed as output view data.

Table 52. Issystemstats attribute values

Attribute	Value
stat_current	The current value of the statistic field.
stat_list	The system history of the reported statistics.
stat_name	The name of the statistic field.
stat_peak	The peak value of the statistic field in the last five minutes.
stat_peak_time	The time that the peak occurred.
sample_time	The time of the sample occurrence.
stat_value	The statistical value at the epoch interval.

Remember: Filtering is supported on the stat_name field using the concise view.

A system summary invocation example

```
lssystemstats
```

The resulting output:

```
stat_name      stat_current  stat_peak  stat_peak_time
cpu_pc         5             6          111123104304
fc_mb          321           327        111123104129
fc_io          2167          2368       111123103904
sas_mb         438           534        111123104104
sas_io         5784          7738       111123104314
iscsi_mb       0             0          111123104359
iscsi_io       0             0          111123104359
write_cache_pc 0             0          111123104359
total_cache_pc 0             0          111123104359
vdisk_mb       321           326        111123104129
vdisk_io       2070          2276       111123103904
vdisk_ms       34            52         111123103954
mdisk_mb       320           329        111123104029
mdisk_io       3135          3340       111123103904
mdisk_ms       15            24         111123104314
drive_mb       440           534        111123104104
drive_io       5765          6572       111123104104
drive_ms       14            21         111123104314
vdisk_r_mb     174           178        111123104324
vdisk_r_io     1064          1180       111123103904
vdisk_r_ms     31            53         111123103954
vdisk_w_mb     146           159        111123104129
vdisk_w_io     1006          1160       111123104129
vdisk_w_ms     38            54         111123104314
mdisk_r_mb     172           177        111123104259
mdisk_r_io     2054          2184       111123103904
mdisk_r_ms     11            18         111123103954
mdisk_w_mb     146           160        111123104129
mdisk_w_io     1081          1229       111123104129
mdisk_w_ms     25            38         111123104314
drive_r_mb     207           356        111123104329
```

```

drive_r_io      2940      3952      111123104104
drive_r_ms      11       18       111123104314
drive_w_mb      231      250      111123104129
drive_w_io      2825     3156     111123104129
drive_w_ms      16       24       111123104314
iplink_mb       0        1        130711190446
iplink_io       0        10       130711190446
iplink_comp_mb   0        250      151014133723

cloud_up_mb     0        0        161118051715
cloud_up_ms     0        0        161118051715
cloud_down_mb   0        0        161118051715
cloud_down_ms   0        0        161118051715

```

A filtered system summary invocation example

```
lssystemstats -filtervalue stat_name=cpu_pc:stat_name=fc_mb -delim :
```

The resulting output:

```

The filtered system summary output:
stat_name:stat_current:stat_peak:stat_peak_time
cpu_pc:5:7:111123104547
fc_mb:319:339:111123104517

```

A historical view-based system summary invocation example

```
lssystemstats -history fc_io
```

The resulting partial output for the historical system summary example:

```

sample_time  stat_name  stat_value
111123104224 fc_io      2120
111123104229 fc_io      2102
111123104234 fc_io      2041
111123104239 fc_io      2211
111123104244 fc_io      2204
111123104249 fc_io      2046
111123104254 fc_io      1997
111123104259 fc_io      2081
111123104304 fc_io      2123
111123104309 fc_io      2030
111123104314 fc_io      1754
111123104319 fc_io      1640
111123104324 fc_io      1759
111123104329 fc_io      1638
111123104334 fc_io      1804
111123104339 fc_io      2011
111123104344 fc_io      2028
111123104349 fc_io      2171
111123104354 fc_io      2055
111123104359 fc_io      2167
111123104404 fc_io      2140
111123104409 fc_io      2111

```

Table 53 provides the possible values that are applicable to the values that are displayed for the **stat_name** attribute.

Table 53. Stat_name field values

Value	Description
compression_cpu_pc	Displays the percentage of allocated CPU capacity utilized for compression.
cpu_pc	Displays the percentage of allocated CPU capacity utilized for the system.

Table 53. Stat_name field values (continued)

Value	Description
fc_mb	Displays the total number of megabytes transferred per second (MBps) for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.
fc_io	Displays the total input/output (I/O) operations transferred per seconds for Fibre Channel traffic on the system. This value includes host I/O and any bandwidth that is used for communication within the system.
sas_mb	Displays the total number of megabytes transferred per second (MBps) for serial-attached SCSI (SAS) traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
sas_io	Displays the total I/O operations transferred per second for SAS traffic on the system. This value includes host I/O and bandwidth that is used for background RAID activity.
iscsi_mb	Displays the total number of megabytes transferred per second (MBps) for iSCSI traffic on the system.
iscsi_io	Displays the total I/O operations transferred per second for iSCSI traffic on the system.
write_cache_pc	Displays the percentage of the write cache usage for the node.
total_cache_pc	Displays the total percentage for both the write and read cache usage for the node.
vdisk_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to volumes during the sample period. Note: Only the write operation value is displayed.
vdisk_io	Displays the average amount of I/O operations transferred per second for read and write operations to volumes during the sample period.
vdisk_ms	Displays the average amount of time in milliseconds that the system takes to respond to read and write requests to volumes over the sample period.
mdisk_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to MDisks during the sample period.
mdisk_io	Displays the average amount of I/O operations transferred per second for read and write operations to MDisks during the sample period.
mdisk_ms	Displays the average amount of time in milliseconds that the system takes to respond to read and write requests to MDisks over the sample period.
drive_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to drives during the sample period
drive_io	Displays the average amount of I/O operations transferred per second for read and write operations to drives during the sample period.
drive_ms	Displays the average amount of time in milliseconds that the system takes to respond to read and write requests to drives over the sample period.
vdisk_w_mb	Displays the average number of megabytes transferred per second (MBps) for read and write operations to volumes during the sample period.
vdisk_w_io	Displays the average amount of I/O operations transferred per second for write operations to volumes during the sample period.
vdisk_w_ms	Displays the average amount of time in milliseconds that the system takes to respond to write requests to volumes over the sample period.
mdisk_w_mb	Displays the average number of megabytes transferred per second (MBps) for write operations to MDisks during the sample period.
mdisk_w_io	Displays the average amount of I/O operations transferred per second for write operations to MDisks during the sample period.


Table 53. Stat_name field values (continued)

Value	Description
mdisk_w_ms	Displays the average amount of time in milliseconds that the system takes to respond to write requests to MDisks over the sample period.
drive_w_mb	Displays the average number of megabytes transferred per second (MBps) for write operations to drives during the sample period
drive_w_io	Displays the average amount of I/O operations transferred per second for write operations to drives during the sample period.
drive_w_ms	Displays the average amount of time in milliseconds that the system takes to respond to write requests to drives over the sample period.
vdisk_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to volumes during the sample period.
vdisk_r_io	Displays the average amount of I/O operations transferred per second for read operations to volumes during the sample period.
vdisk_r_ms	Displays the average amount of time in milliseconds that the system takes to respond to read requests to volumes over the sample period.
mdisk_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to MDisks during the sample period.
mdisk_r_io	Displays the average amount of I/O operations transferred per second for read operations to MDisks during the sample period.
mdisk_r_ms	Displays the average amount of time in milliseconds that the system takes to respond to read requests to MDisks over the sample period.
drive_r_mb	Displays the average number of megabytes transferred per second (MBps) for read operations to drives during the sample period
drive_r_io	Displays the average amount of I/O operations transferred per second for read operations to drives during the sample period.
drive_r_ms	Displays the average amount of time in milliseconds that the system takes to respond to read requests to drives over the sample period.
iplink_mb	Displays the average number of megabytes requested to be transferred per second (MBps) over the IP partnership link during the sample period. This value is calculated before any compression of the data takes place. This value does not include iSCSI host input/output (I/O) operations.
iplink_comp_mb	Displays the average number of compressed megabytes transferred per second (MBps) over the IP Replication link during the sample period. This value is calculated after any compression of the data takes place. This value does not include iSCSI host I/O operations. Note: If compression is disabled, the iplink_mb stats ID value is displayed instead.
cloud_up_mb	Displays the average number of megabytes transferred per second (Mbps) for upload operations to a cloud account during the sample period.
cloud_up_ms	Displays the average amount of time (in milliseconds) it takes for the system to respond to upload requests to a cloud account during the sample period.
cloud_down_mb	Displays the average number of Mbps for download operations to a cloud account during the sample period.
cloud_down_ms	Displays the average amount of time (in milliseconds) it takes for the system to respond to download requests to a cloud account during the sample period.

Istargetportfc

Use the **lstargetportfc** command to generate the lists of worldwide port names (WWPNs) required to set up Fibre Channel (FC) zoning and to display the current failover status of host I/O ports.

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 wildcard characters when using the 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 (").

-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 no data is 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 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 1-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 **ltargetportfc** command:

- port_id
- owning_node_id
- current_node_id
- host_io_permitted
- virtualized

Description

This command generates lists of worldwide port names (WWPNs) required to set up Fibre Channel (FC) zoning. This command also displays the current failover status of host I/O ports.

Table 54 provides the attribute values that can be displayed as output view data.

Table 54. *ltargetportfc* output

Attribute	Description
id	Indicates the ID of the port.

Table 54. *lstorageportfc* output (continued)

Attribute	Description
WWPN	Indicates the WWPN of the port. The value is hexadecimal.
WWNN	Indicates the worldwide node name (WWNN) of the port. The value is hexadecimal.
port_id	Indicates the system port ID. The value is the same as the lstorageportfc port_id field.
owning_node_id	Indicates the ID of the node that owns the port. Note: This node can be offline whether the port is online or offline.
current_node_id	Indicates the ID of the node on which this port is active. The value is blank if the port is not active on any node
nportid	Indicates the nportid hexadecimal value.
host_io_permitted	Indicates whether host I/O operations can run on the port. The values are yes and no.
virtualized	Indicates whether it is a virtualized port. The values are yes and no (which indicates that this port cannot be online on any node other than the owning node).

An invocation example

This example shows a single I/O group with two nodes. One 2-port FC card is installed on each node, and the I/O group's `fctargetportmode` setting is set to disabled.

```
lstorageportfc
```

The detailed resulting output:

id	WWPN	WWNN	port_id	owning_node_id	current_node_id	nportid	host_io_permitted	virtualized
1	500507680140BADD	500507680100BADD	1	1	1	0E2411	yes	no
2	500507680141BADD	500507680100BADD	1	1		000000	no	yes
3	500507680130BADD	500507680100BADD	2	1	1	0E2412	yes	no
4	500507680131BADD	500507680100BADD	2	1		000000	no	yes
5	500507680140BADE	500507680100BADE	1	2	2	0E2413	yes	no
6	500507680141BADE	500507680100BADE	1	2		000000	no	yes
7	500507680130BADE	500507680100BADE	2	2	2	0E2414	yes	no
8	500507680131BADE	500507680100BADE	2	2		000000	no	yes

An invocation example

This example shows a single I/O group with two nodes. One 2-port FC card is installed on each node, and the I/O group's `fctargetportmode` setting is set to transitional.

```
lstorageportfc
```

The detailed resulting output:

id	WWPN	WWNN	port_id	owning_node_id	current_node_id	nportid	host_io_permitted	virtualized
1	500507680140BADD	500507680100BADD	1	1	1	0E2411	yes	no
2	500507680141BADD	500507680100BADD	1	1	1	0E2412	yes	yes
3	500507680130BADD	500507680100BADD	2	1	1	0E2413	yes	no
4	500507680131BADD	500507680100BADD	2	1	1	0E2414	yes	yes
5	500507680140BADE	500507680100BADE	1	2	2	0E2415	yes	no
6	500507680141BADE	500507680100BADE	1	2	2	0E2416	yes	yes
7	500507680130BADE	500507680100BADE	2	2	2	0E2417	yes	no
8	500507680131BADE	500507680100BADE	2	2	2	0E2418	yes	yes

mkquorumapp

Use the **mkquorumapp** command to generate a Java™ application to use for quorum.

Syntax

```
➤ mkquorumapp — [ -ip_6 ] —————➤
```

Parameters

-ip_6
(Optional) Specifies that the quorum application use IPv6 service addresses to connect to nodes.

Description

This command generates a Java application to use for quorum.

An invocation example

This example creates the file `/dumps/ip_quorum.jar` to use for IP quorum on a IPv4 network.

```
mkquorumapp
```

The detailed resulting output:

No feedback

An invocation example

This example creates the file `/dumps/ip_quorum.jar` to use for IP quorum on a IPv6 network.

```
mkquorumapp -ip_6:
```

The detailed resulting output:

mkthrottle

Use the **mkthrottle** command to create a new throttle object and associate it with an object (such as a volume). You can also create offloaded I/O throttling (which is a single clustered system throttle).

Syntax

```
➤ mkthrottle — — -type — offload | vdisk | host | hostcluster | mdiskgrp —————➤
➤ [ -bandwidth — bandwidth_limit_in_mb ] [ -iops — iops_limit ] —————➤
➤ [ -name — throttle_name ] [ -vdisk [ vdisk_id ] [ vdisk_name ] ] [ -host [ host_id ] [ host_name ] ] —————➤
➤ [ -hostcluster [ hostcluster_id ] [ hostcluster_name ] ] [ -mdiskgrp [ mdiskgrp_id ] [ mdiskgrp_name ] ] —————➤
```

Parameters

-type *offload* | *vdisk* | *host* | *hostcluster* | *mdiskgrp*
(Required) Specifies the type of throttle, either *offload* or *vdisk*.

-bandwidth *bandwidth_limit_in_mb*
(Optional) Specifies the bandwidth in MBps. This must be a numeric value from 0 - 268435456.

-iops *iops_limit*

(Optional) Specifies the I/O operations limit. This must be a numeric value 0 - 33554432.

-name *throttle_name*

(Optional) Specifies the throttling object's name. This value must be an alphanumeric string up to 63 characters long.

-vdisk *vdisk_id* | *vdisk_name*

(Optional) Specifies the volume ID or name of the volume to throttle. The value must be a numeric or alphanumeric string.

Note: This keyword must be specified when you specify *-type vdisk*.

This parameter is mandatory for volume throttling but cannot be used for offload throttling.

-host *host_id* | *host_name*

(Optional) Specifies the host ID or name to throttle.

-hostcluster *hostcluster_id* | *hostcluster_name*

(Optional) Specifies the host cluster ID or name to throttle.

-mdiskgrp *mdiskgrp_id* | *mdiskgrp_name*

(Optional) Specifies the MDisk group (storage pool) or name to throttle. This applies to parent or child storage pools.

Description

This command creates a new throttle object and associates it with an object (such as a volume).

Note:

- A throttle object cannot be defined for a host if it is a part of host cluster that already has a host cluster throttle object defined for it.
- If a host cluster does not have a throttle object defined, its member hosts can have individual host throttles defined.
- The storage pool throttle objects for a child pool and a parent pool work independently of each other
- If a volume has multiple copies then throttling is done for the storage pool serving primary copy. Throttling is not applicable for secondary pools that are part of mirrored volumes or stretched cluster implementations.
-

An invocation example for creating a volume throttle of 10000 IOPs and a bandwidth limit of 500 MBps for volume *vdisk0*

```
mkthrottle -type vdisk -iops 10000 -bandwidth 500 -vdisk vdisk0
```

The detailed resulting output:

No feedback

An invocation example for creating offloaded I/O throttling with a bandwidth limit of 500 MBps

```
mkthrottle -type offload -bandwidth 500
```

The detailed resulting output:

No feedback

```
mkthrottle -type host -bandwidth 100 -host host Win2012SP2
```

No feedback

```
mkthrottle -type hostcluster -iops 3000 -hostcluster 0
```

No feedback

```
mkthrottle -type mdiskgrp -iops 40000 -mdiskgrp 0
```

No feedback

Use the **ping** command to diagnose IP configuration problems by checking whether the specified IP address is accessible from the node on which the command is run using the specified IP address.

```
➤ ping -srcip4 source_ipv4_address destination_ipv4_address
      -srcip6 source_ipv6_address destination_ipv6_address
```

-srcip4 *source_ipv4_address destination_ipv4_address*
(Required if **-srcip6** is not specified) Specifies the IPv4 address that sends the ping packet. The IPv4 address must already be bound to a port on the node on which the command is issued. If you do not specify this parameter you must specify **srcip6**.

-srcip6 *source_ipv6_address destination_ipv6_address*
(Required if **-srcip4** is not specified) Specifies the IPv6 address that sends the ping packet. The IPv6 address must already be bound to a port on the node on which the command is issued. If you do not specify this parameter you must specify **srcip4**.

This command checks whether the specified IP address is accessible from the node on which the command is run using the specified IP address.

```
ping -srcip4 192.168.1.51 192.168.1.30
```

The resulting output:

```
PING 192.168.1.51 (192.168.1.51)PING 9.20.136.11 (9.20.136.11) 56(84) bytes of data.
64 bytes from 192.168.1.51: icmp_seq=1 ttl=249 time=0.690 ms
64 bytes from 192.168.1.51: icmp_seq=2 ttl=249 time=0.382 ms
64 bytes from 192.168.1.51: icmp_seq=3 ttl=249 time=0.311 ms
```

```
PING 192.168.1.30 (192.168.1.30)PING 9.20.136.11 (9.20.136.11) 56(84) bytes of data.
64 bytes from 192.168.1.30: icmp_seq=1 ttl=249 time=0.690 ms
64 bytes from 192.168.1.30: icmp_seq=2 ttl=249 time=0.382 ms
64 bytes from 192.168.1.30: icmp_seq=3 ttl=249 time=0.311 ms
```

rmiscsistorageport

Use the **rmiscsistorageport** command to remove established Internet Small Computer Systems Interface (iSCSI) sessions between system nodes and backend iSCSI target.

Syntax

```
➤— rmiscsistorageport — — lsiscsistorageport_row_id —————➤
```

Parameters

lsiscsistorageport_row_id

(Required) Specifies the row ID of the selected row in the output of **lsiscsistorageport** command.

Description

Use this command to remove path groups that are established after specifying **addiscsistorageport** (not individual paths).

Any dependencies must be removed before specifying this command. The ID of the session that is listed (after specifying **lsiscsistorageport**) is used to identify the sessions that must be removed.

A detailed invocation example

After you specify **addiscsistorageport** for discovery and **lsiscsistorageport** to show any added session. Then specify **rmiscsistorageport** to remove sessions that are indicated by view ID 0. No *tgt_user_name* or *target_chap* is required for discovery or session establishment.

```
rmiscsistorageport 0
```

The detailed resulting output:

No feedback

rmnode (SVC) / rmnodecanister (Storwize family products)

The **rmnode** / **rmnodecanister** command deletes a node from the clustered system. You can enter this command any time after a clustered system has been created.

Syntax

```
➤— rmnode — | rmnodecanister — ———— ———— ————➤
                               [ -force ]   [ object_id ]
                               [ object_name ]
```

Parameters

-force

(Optional) Overrides the checks that this command runs. The parameter overrides the following two checks:

- If the command results in volumes 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 or node canister 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 volumes going offline, you force the node or node canister removal and run the risk of losing data from the write cache. The **force** parameter should always be used with caution.

object_id | object_name

(Required) Specifies the object name or ID that you want to modify. The variable that follows the parameter is either:

- The object name that you assigned when you added the node to the clustered system
- The object ID that is assigned to the node (not the worldwide node name)

Description

This command removes a node or node canister from the clustered system. This makes the node or node canister a candidate to be added back into this clustered system or into another system. After the node or node canister is deleted, the other node in the I/O group enters write-through mode until another node or node canister is added back into the I/O group.

Attention: When you run the **rmnode** command to remove the configured hardware for a node:

- Small Computer System Interface-3 (SCSI-3) reservations (through that node) are removed
- Small Computer System Interface-3 (SCSI-3) registrations (through that node) are removed

By default, the **rmnode** / **rmnodecanister** command flushes the cache on the specified node before the node or node canister 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 or node canister with the cache data.

The cache is flushed before the node or node canister is deleted to prevent data loss if a failure occurs on the other node or node canister in the I/O group.

To take the specified node or node canister offline immediately without flushing the cache or ensuring data loss does not occur, run the **rmnode** / **rmnodecanister** command with the **-force** parameter.

Prerequisites:

Before you issue the **rmnode** / **rmnodecanister** command, perform the following tasks and read the following **Attention** notices to avoid losing access to data:

1. Determine which virtual disks (VDIs, or volumes) are still assigned to this I/O group by issuing the following command. The command requests a filtered view of the volumes, where the filter attribute is the I/O group.

```
lsvdisk -filtervalue IO_group_name=name
```

where *name* is the name of the I/O group.

2. Determine the hosts that the volumes are mapped to by issuing the **lsvdiskhostmap** command.

3. Determine if any of the volumes 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 volumes, go to step 5.
 - If you do want to maintain access to some or all of the volumes, 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 or node canister:
 - If this is the last node or node canister in the clustered system, you do not need to turn the power off to the node or node canister. Go to step 5.
 - If this is not the last node or node canister in the cluster, turn the power off to the node or node canister 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 or node canister request.
5. Update the SDD configuration for each virtual path (vpath) that is presented by the volumes that you intend to remove. Updating the SDD configuration removes the vpaths from the volumes. 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 or node canister 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 clustered system. Before you delete the last node or node canister in the clustered system, ensure that you want to destroy the clustered system.
2. If you are removing a single node or node canister and the remaining node or node canister in the I/O group is online, the data can be exposed to a single point of failure if the remaining node or node canister fails.
3. This command might take some time to complete since the cache in the I/O group for that node or node canister is flushed before the node or node canister is removed. If the **-force** parameter is used, the cache is not flushed and the command completes more quickly. However, if the deleted node or node canister is the last node or node canister in the I/O group, using the **-force** option results in the write cache for that node or node canister being discarded rather than flushed, and data loss can occur. The **-force** option should be used with caution.
4. If both nodes or node canisters in the I/O group are online and the volumes are already degraded before deleting the node or node canister, redundancy to the volumes 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 or node canister, the **rmnode / rmnodecanister** command causes the configuration node or node canister to move to a different node or node canister within the clustered system. This process might take a short time: typically less than a minute. The clustered system IP address remains unchanged, but any SSH client attached to the configuration node or node canister might need to reestablish a connection. The management GUI reattaches to the new configuration node or node canister transparently.
2. If this is the last node or node canister in the clustered system or if it is currently assigned as the configuration node, all connections to the system are lost. The user interface and any open CLI sessions are lost if the last node or node canister in the clustered system is deleted. A time-out might occur if a command cannot be completed before the node or node canister is deleted.

An invocation example for rmnode

```
rmnode 1
```

The resulting output:

No feedback

An invocation example for rmnodecanister

```
rmnodecanister 1
```

The resulting output:

No feedback

rmportip

Use the **rmportip** command to remove an Internet Small Computer System Interface (iSCSI) Internet Protocol (IP) address from a node Ethernet port.

Syntax

```
➤➤ rmportip — [ -failover ] — -ip_6 — -node [ node_name | node_id ] — port_id — ➤➤
```

Parameters

-failover

(Optional) Specifies that the failover IP address information be removed for the specified port.

-ip_6

(Required) Specifies that the Internet Protocol Version 6 (IPv6) address be removed for the specified port. If this parameter is not used, the Internet Protocol Version 4 (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, 2, 3, or 4) to apply changes to.

Description

This command removes an IPv4 or IPv6 address from an Ethernet port of a node.

Before unconfiguring an IP from a source Ethernet port, the system checks if there are any sessions established from the selected port to any backend iSCSI controller. You must remove sessions to the backend iSCSI controller using the command before attempting to unconfigure the port.

An invocation example for IPv4

```
rmportip -node 1 1
```

The resulting output:

No feedback

An invocation example for IPv6

```
rmportip -node 1 -ip_6 2
```

The resulting output:

No feedback

rmthrottle

Use the **rmthrottle** command to remove the throttle object associated with any volume.

Syntax

```
►►— rmthrottle — — throttle_id —————►◄  
                  └─ throttle_name ─┘
```

Parameters

throttle_id | *throttle_name*

(Required) Specifies the throttle object ID or name. The value must be a numeric or alphanumeric string up to 63 characters long.

Description

This command removes the throttle object associated with a specified volume.

An invocation example to remove a throttle object with the ID 2

```
rmthrottle 2
```

The detailed resulting output:

No feedback

An invocation example to remove a throttle object with name throttle_vdisk2

```
rmthrottle throttle_vdisk2
```

The detailed resulting output:

No feedback

setclustertime (Discontinued)

Attention: The **setclustertime** command has been discontinued. Use the **setsystemtime** command instead.

setsystemtime

Use the **setsystemtime** command to set the time for the clustered system (system).

Syntax

```
►►— setsystemtime — — -time — time_value —————►◄
```

Parameters

-time *time_value*

(Required) Specifies the time to which the system must be set. This must be in the following format (where Mis month, Dis day, H is hour, m is minute, and Yis year):

MMDDHHmmYYYY

Description

This command sets the time for the a system.

An invocation example

```
setsystemtime -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

```
➤➤ 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 system provides an option to reset the system superuser password to the default value. Use the front panel menu system.

This command allows access if the system superuser password is forgotten. If this feature remains enabled, make sure there is adequate physical security to the system hardware.

You can view or change the status of this feature.

An invocation example

```
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 clustered system (system).

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 **lstimezones** command. The time zone can be changed by running the **settimezone** command.

An invocation example

```
showtimezone -delim :
```

The resulting output:

```
id:timezone  
522:UTC
```

startstats

Use the **startstats** command to modify the interval at which per-node statistics for volumes, managed disks (MDisks), and nodes are collected.

Syntax

```
►►— startstats — — -interval — time_in_minutes —————►►
```

Parameters

-interval *time_in_minutes*

(Required) 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, volumes 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 **lsdumps** 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 volumes, 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 volume 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 volume 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

```
startstats -interval 25
```

The resulting output:

No feedback

stopstats (Deprecated)

The **stopstats** command is deprecated. You can no longer disable statistics collection.


stopcluster (Discontinued)

Attention: The **stopcluster** command has been discontinued. Use the **stopsystem** command instead.

stopsystem

Use the **stopsystem** command to shut down a single node or the entire clustered system in a controlled manner. When you issue this command, you are prompted with a confirmation of intent to process the command.

Syntax



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 volumes 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 volumes going offline, you force the node to shut down, even if it is the last online node in the I/O group.

Remember: 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 system.
- 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 system is shut down.

Description

Use this command to shut down a single node or the entire clustered system in a controlled manner. You are prompted with a confirmation of intent (to process the command) when you specify this command.

If you enter this command with no parameters, the entire system is shut down. All data is flushed to disk before the power is removed.

If you enter this command with either a node ID or node name, the specified node 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 system.

Entering *y* or *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.

If you need to shut down the entire system or a single node, use this command instead of using the power button on the nodes or powering off the main power supplies to the system.

Attention: Do not power off the uninterruptible power supply or remove the power cable from the node.

Storwize V7000: If you need to shut down the system or a single node, use this command instead of using the power button on power supplies, or powering off the mains to the system.

Before shutting down a node or system, complete the following requirements:

1. Quiesce all I/O operations that are destined for this node or system. If you do not quiesce these, failed I/O operations might be reported to your host operating systems.
2. Stop all FlashCopy, Metro Mirror, Global Mirror, and data migration operations.
3. Ensure that all asynchronous deletion operations have completed.

Using this command to shut down a single node fails if shutting down the node makes any volumes inaccessible, or if it is the last node in an I/O group. If you still need to shut down the node, you can use the **-force** option to override these checks.

Important: You can specify `stopssystem -node -reset` to restart the I/O process.

An invocation example

`stopssystem`

The following confirmation prompt is displayed:

Are you sure that you want to continue with the shut down?

Select yes to confirm or no to cancel.

Chapter 8. Clustered system diagnostic and service-aid commands


Clustered system diagnostic and service-aid commands are designed to diagnose and find clustered system (system) 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. The clustered system diagnostic and service-aid commands apply under the service role.

applysoftware

Use the **applysoftware** command to update the clustered system (system) to a new level of system code (code).

Syntax



Parameters

-force

(Optional) Specifies that the update or abort should proceed even if there is a lack of redundancy in the system. Disabling redundancy checking might cause loss of data, or loss of access to data. Use the **force** parameter with the **abort** parameter if one or more nodes are offline.

Important: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information.

-file filename_arg

(Required) Specifies the filename of the installation update package. Copy the update package onto the configuration node before running the **applysoftware** command.

Note: The **file** parameter cannot be used with the **abort** parameter.

-prepare

(Optional) Prepares the system for a manual code level update.

Note: You can:

- Use the **prepare** parameter with the **file** parameter
- Not use the **prepare** parameter with the **abort** parameter
- Not use the **force** parameter with the **prepare** parameter to go to prepared status

-abort

(Required for stopping an update) Specifies that a stalled or prepared update should be stopped, returning the system to the original code level.

Note: The **abort** parameter can be used with the **force** parameter, but not the **file** or **prepare** parameters.

The **abort** parameter can also be used when the **lsupdate** command reports a status of:

- `prepare_failed`
- `prepared` (if all nodes are online)

-complete

(Optional) Start the update completion process. This is either automatic if the `paced` option is not set or `paced` if the `paced` option is set.

-resume

(Optional) Resumes a stalled automatic update, update cancel, or update completion process by retrying the step that stalled. The update completion process becomes `paced` (if that option is set).

-paced

(Optional) Specifies that the update completion should be `paced`. The system does not automatically take any nodes offline and waits for the user to use **-pacednext**.

-pacednext

(Optional) Updates the next node that is part of a `paced` update.

Description

This command starts the update process of the system to a new level of code. The **applysoftware** command applies a level of code to the node as a service action (`Paced update`) to update the specific node, or as an automatic update process that update all of the nodes on a system.

The **applysoftware** command cannot be used in service state, which means the system must be running in order for the command to be used and be successful. This command is synchronous and therefore reports success or failure.

The code package as specified by the file name must first be copied onto the current configuration node in the `/home/admin/update` directory; use the PuTTY secure copy (`scp`) application to copy the file.

If the **applysoftware** command is successful, the **lsupdate** command reports the status is `prepared`. If the **applysoftware** command fails, the **lsupdate** command reports the status is `inactive`.

If specified, the **prepare** parameter must succeed in order to successfully update. It is recommended to use the same package for the `prepare` as the actual update. The **prepare** parameter can be canceled by using the `abort` parameter (even after the system is `prepared`) as long as the **lsupdate** command reports the status as `prepared`.

Important: The **prepare** parameter might time out. If this occurs, the **prepare** parameter causes an asynchronous condition, and the **lsupdate** command reports the `prepare` as `"preparing"`. If this occurs then wait until **lsupdate** reports the update as `prepared` before proceeding with the manual update process.

The command completes as soon as the update process is successful. The command fails and the update package is deleted if:

- The given package fails an integrity check due to corruption.
- Any node in the system has a hardware type not supported by the new code.
- The new code level does not support updates from the currently installed code.

- The code level of a remote system is incompatible with the new code.
- There are any volumes 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 update. Before proceeding, use the **lsdependentvdisks** command with the **node** parameter to list the node-dependent volumes at the time the command is run. If the command returns an error, move the quorum disks to MDiskS that are accessible through all nodes. Rerun the command until no errors are returned.

The actual update completes asynchronously.

An invocation example

```
applysoftware -file filename_arg
```

The resulting output:

No feedback

An invocation example

```
applysoftware -prepare -file INSTALL_6.4.0.0
```

The resulting output:

No feedback

An invocation example

```
applysoftware -abort
```

The resulting output:

No feedback

An invocation example

```
applysoftware -file softwareupdate
```

The resulting output:

No feedback

An invocation example

```
applysoftware -complete -force
```

The resulting output:

No feedback

An invocation example

```
applysoftware -resume -paced
```

The resulting output:

No feedback

An invocation example

```
applysoftware -pacednext -force
```

The resulting output:

No feedback

caterlog (Deprecated)

The **caterlog** command has been deprecated. Use the **lseventlog** command instead.

caterlogbyseqnum (Deprecated)

The **caterlogbyseqnum** command has been deprecated. Use the **lseventlog** command instead.

cherrstate (Deprecated)

The **cherrstate** command has been deprecated. Use the **cheventlog** command instead.

chdnsserver

Use the **chdnsserver** command to change a Domain Name System (DNS) server Internet Protocol (IP) address or name on a clustered system (system).

Syntax

```
➤ chdnsserver — [ -ip — ip_address ] [ -name — dns_name ] [ dns_name | dns_id ] ➤
```

Parameters

-ip *ip_address*

(Optional) Specifies the DNS server IP address. The value must be in standard IPv4 or IPv6 format.

-name *DNS_name*

(Optional) Specifies a unique name for the system DNS server that is being changed.

dns_name | *dns_id*

(Required) Specifies DNS unique name or ID associated with the DNS server that is being changed.

Description

This command changes DNS server parameters (such as an IP address or name).

An invocation example

```
chdnsserver -ip 192.168.48.220 1
```

The resulting output:

No feedback

An invocation example

```
chdnsserver -name dns1 1
```

The resulting output:

No feedback

cheventlog

Use the **cheventlog** command to modify events in the event log.

Syntax

```
►► cheventlog --fix--sequence_number_____◄◄
```

Parameters

-fix *sequence_number*
(Required) Mark an unfixed event as fixed.

Description

Important: You must specify the **-fix** parameter.

An invocation example to mark an event fixed

```
cheventlog -fix 120
```

The resulting output:

No feedback

chsyslogserver

Use the **chsyslogserver** command to modify the parameters of an existing syslog server.

Syntax

```
►► chsyslogserver — [ -name — server_name ] [ -ip — ip_address ] _____◄◄  
  
► [ -facility — facility ] [ -error [ on ] [ off ] ] [ -warning [ on ] [ off ] ] _____◄◄  
  
► [ -info [ on ] [ off ] ] [ -cadf [ on ] [ off ] ] [ syslog_server_name  
  syslog_server_id ] _____◄◄
```

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. This parameter is mutually-exclusive with **-cadf**.

-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.

-cadf on | off

(Optional) Specifies that Cloud Auditing Data Federation (CADF) data reporting be turned on or off. Any syslog notifications sent to the server are formatted to the CADF standard. This parameter is mutually-exclusive with **-facility**.

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 **lssyslogserver** command to obtain this ID.

If you disable CADF notifications for a syslog server that has CADF notification enabled, the facility value must be set to 0.

An invocation example

```
chsyslogserver -facility 5 2
```

The resulting output:

No feedback

An invocation example

```
chsyslogserver -cadf on 0
```

The resulting output:

No feedback

clearerrlog

Use the **clearerrlog** command to clear all entries from the event log including status events and any unfixed errors.

Syntax

```
➤ 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 event log. The entries are cleared even if there are unfixed events in the log. It also clears any status events that are in the log.

Attention: This command is destructive. Use it only when you have either rebuilt the clustered system or have fixed a major problem that has caused entries in the event log that you do not want to manually fix.

An invocation example

```
clearerrlog -force
```

The resulting output:

No feedback

cpfabricdumps (Discontinued)

The **cpfabricdumps** command is discontinued. Use the **cpdumps** command instead.

dumperrlog

Use the **dumperrlog** command to dump the contents of the event log to a text file.

Syntax

```
►► 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 clustered system (system) event 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 event-log dump files are kept on the system. When the 11th dump is made, the oldest existing dump file is overwritten.

Event log dump files are written to `/dumps/elog`s. The contents of this directory can be viewed using the **lsdumps** command.

Files are not deleted from other nodes until you issue the **cleardumps** command.

Note: The DMP family is printed at the end of all events in the event log.

An invocation example

```
dumperrlog -prefix testerrorlog
```

The resulting output:

No feedback

finderr

Use the **finderr** command to analyze the event log for the highest severity unfixed event.

Syntax

►— finderr —◄

Parameters

None

Description

The command scans the event log for any unfixed events. Given a priority ordering within the code, the highest priority unfixed event is returned to standard output.

You can use this command to determine the order in which to fix the logged event.

An invocation example

```
finderr
```

The resulting output

Highest priority unfixed event code is [1010]

setevent (Discontinued)

Attention: The **setevent** command is discontinued. SNMP notification can be configured using the following commands: **svctask mksnmpserver**, **svctask chsnmpserver**, **svctask rmsnmpserver**, and **svcinfo lssnmpserver**.

lscimomdumps (Deprecated)

The **lscimomdumps** command is deprecated. Use the **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

►► `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 clustered system at a time.

An invocation example

```
lscopystatus
```

The resulting output:

```
status
active
```

Isdumps

Use the **`lsdumps`** command to display a list of files in a particular dumps directory on one of the nodes in the clustered system (system).

Syntax

►► `lsdumps` — `-nohdr` — `-delim — delimiter` — `-prefix — directory_name` —

► — `node_name` — `node_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, 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/cloud
- /dumps/easytier
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/mdisk
- /home/admin/update
- /dumps/drive
- /dumps/enclosure

node_name | *node_id*

(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.

Use the **lsdumps** 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. If you do not specify a node, the available dumps on the configuration node are listed.

An invocation example to list the files in /dumps on the configuration node

```
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 files in /dumps/easytier on the configuration node

```
lsdumps -prefix /dumps/easytier/ node_1
```

The resulting output:

```
id filename
0 dpa_heat.78RE5LV-1.150705.074636.data
1 dpa_log_78RE5LV-1_20150707062320_00000000.xml.gz
```

Isdnsserver

Use the **lsdnsserver** command to list information for any Domain Name System (DNS) servers in the clustered system (system).

Syntax

```
lsdnsserver — [ -delim — delimiter ] [ -nohdr ] [ dns_name — dns_id ]
```

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.

-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.

dns_name | *dns_id*

(Optional) Specifies the DNS server name or ID for which to display details. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command lists information for any DNS servers in the system.

This value provides the attribute values that can be displayed as output view data.

Table 55. *Isdnsserver* output

Attribute	Description
id	Indicates the DNS server ID. The value must be a number.
name	Indicates the DNS server name. The value must be an alphanumeric string.
type	Indicates the DNS server Internet Protocol (IP) address type. The value must be a standard IPv4 or IPv6 address.

Table 55. *lscnsserver* output (continued)

Attribute	Description
IP_address	Indicates the IP address of the DNS server. The value must be a standard IPv4 or IPv6 address.

A concise invocation example

```
lscnsserver
```

The resulting output:

id	name	type	IP_address
0	DNS1	ipv6	2801:0000:0000:0000:0000:0000:0100
1	DNS2	ipv4	192.168.44.34

A detailed invocation example

```
lscnsserver 1
```

The resulting output:

```
id 1
name DNS2
type ipv4
IP_address 192.168.44.34
```

lserrlogbyfcconsistgrp (Deprecated)

The **lserrlogbyfcconsistgrp** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyfcmap (Deprecated)

The **lserrlogbyfcmap** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyhost (Deprecated)

The **lserrlogbyhost** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyiogrp (Deprecated)

The **lserrlogbyiogrp** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbymdisk (Deprecated)

The **lserrlogbymdisk** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbymdiskgrp (Deprecated)

The **lserrlogbymdiskgrp** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbynode (Deprecated)

The **lserrlogbynode** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyrconsistgrp (Deprecated)

Attention: The **lserrlogbyrconsistgrp** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyrcrelationship (Deprecated)

The **lserrlogbyrcrelationship** command has been deprecated. Use the **lseventlog** command instead.

lserrlogbyvdisk (Deprecated)

The **svcinfo lserrlogbyvdisk** command has been deprecated. Use the **svcinfo lseventlog** command instead.

lserrlogdumps (Deprecated)

The **lserrlogdumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

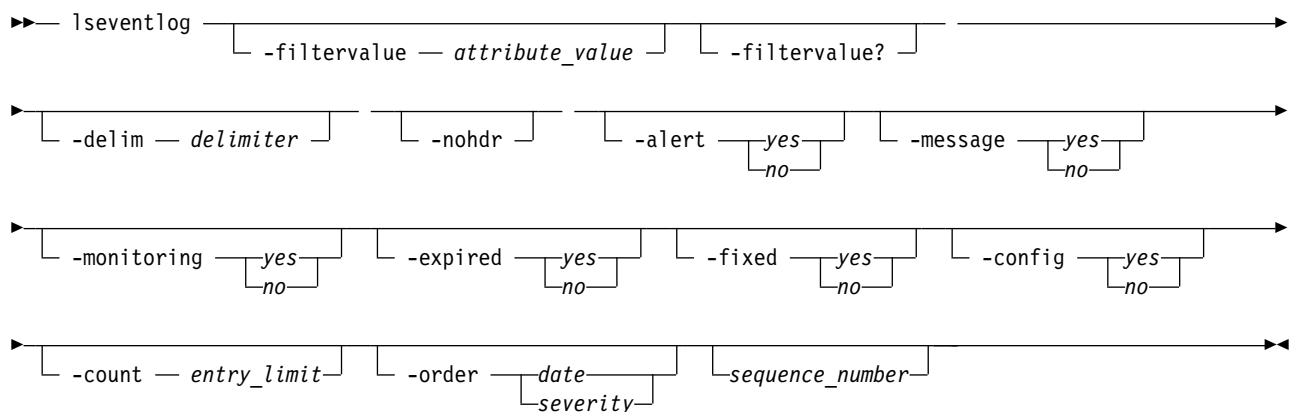
lsfeaturedumps (Deprecated)

The **lsfeaturedumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lseventlog

Use the **lseventlog** command to display a concise view of the system event log, or a detailed view of one entry from the log.

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 when using the 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:

```
lseventlog -filtervalue "object_name=ob*"
```

-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.

-alert yes | no

(Optional) Includes (or excludes) events with *alert* status.

-message yes | no

(Optional) Includes events with *message* status.

-monitoring yes | no

(Optional) Includes events with *monitoring* status.

-expired yes | no

(Optional) Includes (or excludes) events with *expired* status.

-fixed yes | no

(Optional) Includes (or excludes) events with *fixed* status.

-config yes | no

(Optional) Includes (or excludes) events with *config* status.

-count entry_limit

(Optional) Indicates the maximum number of events to display.

-order date | severity

(Optional) 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. Fixed alerts and messages

-filtervalue?

(Optional) Displays a list of valid filter attributes for the **-filtervalue attribute=value** parameter:

- copy_id
- error_code
- event_count
- event_id
- fixed
- last_timestamp
- object_id

- object_name
- object_type
- report_node
- reporting_node_name
- root_sequence_number
- sequence_number status

sequence_number

(Optional) 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.

The default values for included events are:

- alert=yes
- config=no
- expired=no
- fixed=no
- message=yes
- monitoring=no

Table 56 provides the attribute values that can be displayed as output view data.

Table 56. lseventlog output

Attribute	Description	Value
machine_type	Node machine type and model number	The value is an alphanumeric string up to 7 characters long.
serial number	Node serial number	The value is an alphanumeric string up to 7 characters long.
sequence_number	Sequence number of the event	The value is numeric from 0 to 8000000.
first_timestamp	When the event was added to the log	The value is in the format YYMMDDHHMMSS.
first_timestamp_epoch	When the event was added to the log (in seconds) after the epoch occurs	This is a numeric 32-bit value.
last_timestamp	When the event was most recently updated	The value is in the format YYMMDDHHMMSS.
last_timestamp_epoch	Most recent update (in seconds) after an epoch for an event	This is a numeric 32-bit value.
fixed_timestamp	Time stamp when event is fixed	The value is in the format YYMMDDHHMMSS.
fixed_timestamp_epoch	Time stamp (in seconds) when an event is fixed after an epoch occurs	The value is a numeric string.
fru	Field-replaceable unit (FRU) for error or event; this field contains probable FRUs (separated by commas)	The value is an ASCII string up to 255 characters long.

Table 56. lseventlog output (continued)

Attribute	Description	Value
object_type	The type of the object the event is logged against	<p>The values are:</p> <ul style="list-style-type: none"> • mdisk • mdiskgrp • volume (or volume copy) • node • host • io_grp (iogroup in dumperrlog) • fc_consist_grp (fcgrp in dumperrlog) • rc_consist_grp(rcgrp in dumperrlog) • fc_map (fcmap in dumperrlog; <i>flash</i> in catterrlog) • rc_relationship (rcmap in dumperrlog; <i>remote</i> in catterrlog) • cluster • controller (device in catterrlog and dumperrlog) • quorum • migrate • email_server (<i>emailserver</i> in catterrlog and dumperrlog) • enclosure • drive
object_id	ID of the object the event is logged against	<p>This is a numeric 64-bit value. It is displayed in decimal for all object types other than clustered systems.</p> <p>For a clustered system, this value is hexadecimal but is blank for events with cluster object types.</p>
object_name	Name of the object the event is logged against.	This value is the object name format and is blank if the object was deleted or does not have a name
copy_id	Volume copy ID the event is logged against	This is a numeric value 0 to 1; it is blank if not a vdiskcopy event
reporting_node_id	ID of the node that reported the event	This is a numeric 64-bit value that is blank if the event is reported by the clustered system
reporting_node_name	Name of the node that reported the event.	This value is the object name format and is blank if node is deleted or event is reported by the clustered system
root_sequence_number	<p>Sequence number of the root or causal event</p> <p>Important: If the event is directly caused by another event then the <code>sequence_number</code> of the related event is shown here.</p>	This is a numeric value from 1 to 8000000; blank if there is no root or if the event is not directly caused by another event

Table 56. Iseventlog output (continued)

Attribute	Description	Value
event_count	Number of reported events that have been combined into this event	This is a numeric 32-bit value.
status	Event category	The values are: <ul style="list-style-type: none"> • alert • message • monitoring • expired
fixed	Indicates whether the event was marked fixed (for an alert) or read (for a message)	The values are: <ul style="list-style-type: none"> • yes • no (for events that cannot be fixed, or are not fixed)
auto_fixed	Indicates if event is marked fixed by the code	The values are: <ul style="list-style-type: none"> • yes • no (for events that cannot be fixed, or are not fixed)
notification_type	Type of event notification	The values are: <ul style="list-style-type: none"> • error • warning • informational • none
event_id	Event ID	This is a 6-digit numeric value.
event_id_text	Description associated with the event ID	This is a text value with a maximum of 200 bytes. This appears in CLI requested language.
error_code	Error code associated with this event	This is a 4-digit numeric value but is blank if there is no error code.
error_code_text	Description associated with the error code	This is a text value with a maximum of 200 bytes that is blank if there is no error code. This appears in CLI requested language.
description	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 56. lseventlog output (continued)

Attribute	Description	Value
sense1	Sixteen bytes of hex-encoded sense data; least significant byte is on the left	The value is sixteen two-character hexadecimal numbers separated by spaces.
sense2		
sense3		
sense4		
sense5		
sense6		
sense7		
sense8		

Invocation examples

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 120
```

```
sequence_number 120
first_timestamp 111130100419
first_timestamp_epoch 1322647459
last_timestamp 111130100419
last_timestamp_epoch 1322647459
object_type node
object_id 1
object_name node1
copy_id
reporting_node_id 1
reporting_node_name node1
root_sequence_number
event_count 1
status alert
fixed yes
auto_fixed no
notification_type error
event_id 073003
event_id_text More/Less fibre channel ports operational
error_code 1060
error_code_text Fibre Channel ports not operational
```

```

machine_type 21458F4
serial_number 75BZPMA
fru none
fixed_timestamp 111202141004
fixed_timestamp_epoch 1322835004

```

```

sense1 03 03 00 00 00 00 00 00 00 00 00 00 00 00 00 00
sense2 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
sense3 00 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 00
sense5 00 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 00
sense7 00 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 00

```

Isservicestatus

Use the **Isservicestatus** command to display the current status of a node.

Syntax

```

▶▶ — sainfo — — Isservicestatus —————▶▶
                        |panel_name|

```

Parameters

panel_name

(Optional) 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 fails if the *panel_name* ID is not in the list returned by **Isservicenodes**. This output is returned as the node status on all Universal Serial Bus (USB) flash drive commands.

Note: For 2076 nodes the *panel name* is the value of the enclosure ID and canister ID or the enclosure serial number and canister location.

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 clustered system (system), to obtain the vital product data (VPD) and error status.

Table 57 shows possible outputs.

Table 57. Isservicestatus output

Attribute	Value
panel_name	The front panel name, enclosure IDs, or canister IDs that identify the node.
console_ip	An Internet Protocol (IP) Version 4 or 6 address Note: This field might be blank if the node is not present in a system.
has_nas_key	yes no Note: This field might be blank if the node is not present in a system.
cluster_id	Specifies the ID of a system.
cluster_name	Specifies the name of a system. When you use this parameter, the detailed view of the specific system 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.

Table 57. Isservicestatus output (continued)

Attribute	Value
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	4F2 CF8 CG8 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 version of code for the node.
node_sw_build	The build string version of code on the node.
cluster_sw_build	The CSM build that the system is running.
node_error_count	The number of node errors.
node_error_data	The type of node errors.
FC_port_count	The number of FC ports.
FC_port_id	The beginning of repeating fields for each FC port; the whole set of fields indicated is repeated for each port.
port_status	This should match the port on the front panel, enclosure, or canister.
port_speed	This should match the port speed on the front panel, enclosure, or canister.
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	10Mbps 100Mbps 1Gbps 10Gbps full half
MAC	A single MAC address.
vnport_count	Number of VN ports created on top of each physical FCoE port.
vnport_id	The VN port ID.
vnport_wwpn	The WWPN assigned to the VN port.
vnport_FCF_mac	The MAC address for the FCF to which the VN port is connected.
vnport_vlanid	The VLAN ID used by the VN port. The value is blank for FC ports.
product_mtm	The machine type and model.
product_serial	The node serial number.
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
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.
Battery_count	The number of expected batteries (two).
Battery_id	Specifies the ID of the slot the battery is in.
Battery_status	The status is missing, failed, charging, or active.
Battery_FRU_part	The FRU part number of the battery.
Battery_part_identity	The IIS FRU identity of the battery (includes the serial number).
Battery_fault_led	The fault light-emitting diode (LED) status.
Battery_charging_status	The battery charging status.
Battery_cycle_count	The number of charge or discharge cycles performed by the battery.

Table 57. *Isservicestatus output (continued)*

Attribute	Value
Battery_powered_on_hours	The number of hours the battery is in a powered node.
Battery_last_recondition	A system timestamp indicating the last successful gas gauge calibration.
Battery_midplane_FRU_part	The FRU part number of the battery midplane
Battery_midplane_part_identity	The 11S FRU identity of the battery midplane (which includes the serial number).
Battery_midplane_FW_version	The firmware version running on the battery midplane.
Battery_power_cable_FRU_part	The FRU part number of the battery power cable.
Battery_power_sense_cable_FRU_part	The FRU part number of the battery midplane power sense cable.
Battery_comms_cable_FRU_part	The FRU part number of the battery midplane communication cable.
Battery_EPOW_cable_FRU_part	The FRU part number of the battery midplane EPOW cable.
PSU_count	N/A
PSU_id	N/A
PSU_status	N/A
local_fc_port_mask	Indicates the FC I/O ports that a system can use for node-to-node communications on a local system if those FC I/O ports exist on a node. The value is 64 binary bits.
partner_fc_port_mask	Indicates the FC I/O ports that a system can use for system-to-system communications on a partnered system if those FC I/O ports exist on a node. The value is 64 binary bits.
cluster_topology	Indicates the system topology (set using the chsystem command).
site_id	Indicates the site node value.
site_name	Indicates the site name.
identify_LED	Indicates the node or node canister identify LED state (on, off, or blank).
password_reset_enabled	Indicates if the superuser password reset is enabled (yes or no).

An invocation example

```
sainfo lsservicestatus
```

The resulting output:

```
panel_name 150434
cluster_id 000002006ee1445e
cluster_name Cluster_192.168.8.241
cluster_status Active
cluster_ip_count 2
cluster_port 1
cluster_ip 192.168.8.241
cluster_gw 192.168.8.1
cluster_mask 255.255.255.0
cluster_ip_6
cluster_gw_6
cluster_prefix_6
cluster_port 2
cluster_ip
cluster_gw
cluster_mask
cluster_ip_6
cluster_gw_6
cluster_prefix_6
node_id 1
node_name node1
node_status Active
config_node Yes
hardware CF8

service_IP_address
service_gateway
service_subnet_mask
service_IP_address_6
service_gateway_6
service_prefix_6
service_IP_mode dhcpfallback
node_sw_version 6.4.0.0
node_sw_build 64.8.1205180000
cluster_sw_build 64.8.1205180000
```

```

node_error_count 0
fc_ports 4
port_id 1
port_status Active
port_speed 8Gb
port_WWPN 500507680140a22f
SFP_type Short-wave
port_id 2
port_status Active
port_speed 8Gb
port_WWPN 500507680130a22f
SFP_type Short-wave
port_id 3
port_status Active
port_speed 8Gb
port_WWPN 500507680110a22f
SFP_type Short-wave
port_id 4
port_status Active
port_speed 8Gb
port_WWPN 500507680120a22f
SFP_type Short-wave
ethernet_ports 4
ethernet_port_id 1
port_status Link Online
port_speed 1Gb/s - Full
MAC 00:21:5e:db:30:38
vnport_count 0
ethernet_port_id 2
port_status Not Configured
port_speed
MAC 00:21:5e:db:30:3a
vnport_count 0
ethernet_port_id 3
port_status Not Configured
port_speed 10Gb/s - Full
MAC 00:00:c9:bc:6f:22
vnport_count 0
ethernet_port_id 4
port_status Not Configured
port_speed 10Gb/s - Full
MAC 00:00:c9:bc:6f:20
vnport_count 0
product_mtm 2145-CF8

product_serial 75HAXYA
time_to_charge 0
Battery_charging 0
Battery_count 2
Battery_id 1
Battery_status active
Battery_FRU_part 12Z9876
Battery_part_identity 11S98Z1234YM11BG123456
Battery_fault_led OFF
Battery_charging_Status charged
Battery_cycle_count 5
Battery_powered_on_hours 12345
Battery_last_recondition 130629123456
Battery_id 2
Battery_status failed
Battery_FRU_part 12Z9876
Battery_part_identity 11S98Z1234YM11BG234567
Battery_fault_led ON
Battery_charging_Status charged
Battery_cycle_count 5
Battery_power_on_hours 12345
Battery_last_recondition 130702123400

```


Battery_midplane_FRU_part 12Z9880
Battery_power_cable_FRU_part 12Z9881
Battery_power_sense_cable_FRU_part 12Z9882
Battery_comms_cable_FRU_part 12Z9883
Battery_EPOW_cable_FRU_part 12Z9884

dump_name 150434
node_WWNN 500507680100a22f
disk_WWNN_suffix 0A22F
panel_WWNN_suffix 0A22F
UPS_serial_number
UPS_status
enclosure_WWNN_1
enclosure_WWNN_2
node_part_identity
node_FRU_part
enclosure_identity
PSU_count
PSU_id
PSU_status
PSU_id
PSU_status

node_location_copy
node_product_mtm_copy
node_product_serial_copy
node_WWNN_1_copy
node_WWNN_2_copy
latest_cluster_id
next_cluster_id
console_IP 192.168.8.241:443
has_nas_key no
fc_io_ports 6
fc_io_port_id 1
fc_io_port_WWPN 500507680140a22f
fc_io_port_switch_WWPN 200000051e630f9a
fc_io_port_state Active
fc_io_port_FCF_MAC N/A
fc_io_port_vlanid N/A
fc_io_port_type FC
fc_io_port_type_port_id 1
fc_io_port_id 2
fc_io_port_WWPN 500507680130a22f
fc_io_port_switch_WWPN 200400051e630f9a
fc_io_port_state Active
fc_io_port_FCF_MAC N/A
fc_io_port_vlanid N/A
fc_io_port_type FC
fc_io_port_type_port_id 2
fc_io_port_id 3
fc_io_port_WWPN 500507680110a22f
fc_io_port_switch_WWPN 200000051e7ded49
fc_io_port_state Active
fc_io_port_FCF_MAC N/A
fc_io_port_vlanid N/A
fc_io_port_type FC
fc_io_port_type_port_id 3
fc_io_port_id 4
fc_io_port_WWPN 500507680120a22f
fc_io_port_switch_WWPN 200400051e7ded49
fc_io_port_state Active
fc_io_port_FCF_MAC N/A
fc_io_port_vlanid N/A
fc_io_port_type FC
fc_io_port_type_port_id 4
fc_io_port_id 5
fc_io_port_WWPN 500507680150a22f

Issyslogserver

Syntax



Note: If there is no data to be displayed, headings are not displayed.

(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 clustered system.

This table provides the attribute values that can be displayed as output view data.

Table 58. lssyslogserver output

Attribute	Description
id	Indicates the syslog server ID. The value must be a number.
name	Indicates the syslog server name. The value must be an alphanumeric string.
IP_address	Indicates the syslog server Internet Protocol (IP) address. The value must be a valid IP address.
facility	Indicates the syslog server facility value. The value must be a number 0 - 7 but is blank for a CADF notification-enabled server.
error	Indicates whether error messages are on. The values are on or off.
warning	Indicates whether warning messages are on. The values are on or off.
info	Indicates whether informational messages are on. The values are on or off.
cadf	Indicates whether CADF data reporting is enabled or not for the syslog server. The values are on or off.

A concise invocation example

```
lssyslogserver -delim :
```

The concise resulting output:

```
id:name:IP_address:facility:error:warning:info:cadf
0:syslog0:192.135.60.4::on:on:on:on
1:newserver:192.136.70.7:4:on:off:off:on
```

A detailed invocation example

```
lssyslogserver 0
```

The detailed resulting output:

```
id 0
name syslog0
IP_address 192.135.60.4
facility
error on
warning on
info on
cadf on
```

Issoftwaredumps (Deprecated)

The **lssoftwaredumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

Issoftwareupgradestatus (Deprecated)

The **lssoftwareupgradestatus** command is deprecated. Use the **lsupdate** command instead.

lsupdate

Use the **lsupdate** command to display a system's machine code (code) upgrade status.

Syntax

```
lsupdate [-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.

ne

Description

The following are the upgrade status states:

status Indicates the overall update-related status of the system. The values are:

- success, which indicates that all updating activity is complete.
- system_preparing, which indicates that the system is preparing a manual update.
- system_prepared, which indicates that the system is ready to starting a manual update.
- system_prepare_failed, which indicates that the system could not start a manual update.

Note: Check the event log.

- system_initializing, which indicates that the system is readying nodes for an update.
- system Updating, which indicates that the nodes are being updated.
- system Updating Pausing, which indicates that the system is pausing before continuing to update the nodes.
- system Committing, which indicates that all nodes have been updated and the system is readying the new code.
- system Stalled, which indicates that an update has stalled because of unexpected node problems.
- system Stalled Non-Redundant, which indicates that an update has stalled because of dependent volumes.
- system Restoring, which indicates that a stalled update has been cancelled by the user.

Note: The previous code version is being restored.

- `system_restoring_pausing`, which indicates that the system is pausing before continuing to restore the nodes.
- `system_restoring_stalled_non_redundant`, which indicates that an update has been cancelled and then stalled because of dependent volumes..
- `system_manual_update`, which indicates that a manual update is in progress.
- `system_completion_required`, which indicates that all nodes have updated and management functions are available, but further system changes are necessary to complete the update.

Note: Check the event log.

- `system_completing`, which indicates that an automatic update completion is in progress.
- `system_completing_pausing`, which indicates that automatic update completion is paused.
- `system_completing_paced`, which indicates that a paced update completion is in progress.
- `system_completing_stalled`, which indicates that an automatic update completion has stalled because of an unexpected problem.
- `enclosures`, which indicates that enclosure firmware is being updated.
- `enclosures_stalled`, which indicates that an enclosure firmware update has stalled because of an unexpected problem or the enclosure has a lack of redundancy.

Note: On a Storwize V7000 Gen1 system, battery maintenance can leave an enclosure without redundancy for several hours.

- `drives`, which indicates that drive firmware is being updated.

event_sequence_number

Indicates an event describing any current problem with the code update. The value must be a numeric string in decimal format (or blank).

progress

Indicates the completion percentage of the current update activity in terms of number of objects updated (rather than time elapsed). The value must be a numeric string (decimal) between 0 and 100.

estimated_completion_time

Indicates estimated completion time of current update activity. It is valid only if the current update activity is automatic and has not stalled. The value must be in the format YYMMDDHHMMSS (or blank).

suggested_action

Indicates the actions that help the update progress. The value must be:

- `start`, which indicates that the system is ready for a new update to start. No update is in progress or prepared and all nodes are online.
- `wait`, which indicates that the system is busy (no action is required) because an update is in progress.
- `fix`, which indicates that an update cannot continue because of a problem. Check the event log, specifically the **event_sequence_number** output. Nodes are offline.
- `resume_cancel`, which indicates that the update is stalled but can be resumed or canceled. Nodes are online.
- `resume`, which indicates that the update completion is stalled but can be resumed. Nodes are online.
- `complete`, which indicates the system update is complete and update completion must be issued. Nodes are online.
- `pacednext`, which indicates that the paced update is in progress and the next node must be scheduled for updating. Nodes are online.
- `manual`, which indicates that a manual update is in progress.

system_new_code_level

Indicates that a new level of code is being updated. The value must be the build version (or blank if not updating or restoring the system).

system_forced

Indicates any current node-related activity in forced mode (ignoring dependent volumes). The values are yes or no.

system_next_node_status

Indicates the status of the next node in the current node-related update activity. The values are:

- **waiting**, which indicates that the node is ready for updating and that the system is waiting (typically for multipathing failover).
- **ready**, which indicates that the node is ready for updating, and the update activity is paced. You must start the update manually.
- **updating**, which indicates that the node is updating.
- **stalled**, which indicates that the node is going to be updated next, but the update is stalled.
- **none**, which indicates that there is no node to update.

system_next_node_time

Indicates the time that the next node update will start. It is valid only if the **system_next_node_status** is waiting. The value must be in the format YYYYMMDDHHMMSS x (or blank).

system_next_node_id

Indicates the ID of the next node in the current node-related update. The value must be a numeric string (or blank).

system_next_node_name

Indicates the name of the next node in the current node-related update. The value must be an alphanumeric string (or blank).

An invocation example of an update

```
lupdate
```

The resulting output:

```
status system_updating
event_sequence_number
progress 50
estimated_completion_time 140522093020
suggested_action wait
system_new_code_level 7.4.0.1 (build 99.2.141022001)
system_forced no
system_next_node_status updating
system_next_node_time
system_next_node_id 2
system_next_node_name node2
```

An invocation example of a paced update

```
lupdate
```

The resulting output:

```
status system_completing_paced
event_sequence_number
progress 75
estimated_completion_time
suggested_action pacednext
system_new_code_level
system_forced no
```

```
system_next_node_status ready
system_next_node_time
system_next_node_id 4
system_next_node_name node4
```

mkdnsserver

Use the **mkdnsserver** command to configure a new Domain Name System (DNS) server for a clustered system (system).

Syntax

```
➤— mkdnsserver — — -ip — ip_address — — ———— ➤
                               └ — -name — DNS_name — ┘
```

Parameters

-ip *ip_address*

(Required) Specifies the DNS server Internet Protocol (IP) address. The value must be in standard IPv4 or IPv6 format. Depending on the format you specify the system will validate the format to make sure it is correct.

-name *DNS_name*

(Optional) Specifies a unique name for the system DNS server being created. If a DNS server name is not specified, a unique name is generated and then assigned to the DNS server.

Description

This command configures a new DNS server for a system.

Use the **-ip** parameter to specify the DNS server IP address. You can enter any valid IPv4 or IPv6 address. The system validates the format of specified IP address to make sure that it is correct.

An invocation example

```
mkdnsserver -ip 192.168.44.34
```

The resulting output:

```
DNS Server id [0] successfully created
```

An invocation example

```
mkdnsserver -ip 2801:0000:0000:0000:0000:0000:0100
```

The resulting output:


```
DNS Server id [1] successfully created
```

mksyslogserver

Use the **mksyslogserver** command to create a syslog server to receive notifications.

Syntax

```
➤— mksyslogserver — — ———— ➤
                               └ — -name — server_name — ┘ — -ip — ip_address ———— ➤
```



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 Internet Protocol (IP) address of the syslog server. This must be a valid Internet Protocol Version 4 (IPv4) or Internet Protocol Version 6 (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**.

-cadf **on** | **off**

(Optional) Specifies that Cloud Auditing Data Federation (CADF) data reporting be turned on or off. Any syslog notifications sent to the server are formatted to the CADF standard. This parameter is mutually-exclusive with **-facility**.

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 six syslog servers.

An invocation example

```
mksyslogserver -ip 1.2.3.4
```

The resulting output:

```
Syslog Server id [2] successfully created
```


An invocation example

```
mksyslogserver -ip 9.193.231.37 -error on -warning on -info off -cadf on
```

The resulting output:

```
Syslog Server id [2] successfully created
```

rmdnsserver

Use the **rmdnsserver** command to remove a Domain Name System (DNS) server from a clustered system (system).

Syntax

```
➤— rmdnsserver — [ dns_name | dns_id ] —————➤
```

Parameters

dns_name | *dns_id*

(Required) Specifies the ID or name of the Domain Name System (DNS) server to be removed from the system. The value must be a number for the ID and an alphanumeric string for the name.

Description

This command removes a DNS server from a system.

An invocation example

```
rmdnsserver 1
```

The resulting output:

```
No feedback
```

rmsyslogserver

Use the **rmsyslogserver** command to delete the specified syslog server.

Syntax

```
➤— 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 **lssyslogserver** command to obtain this ID.

An invocation example

```
rmsyslogserver 2
```

The resulting output (if the command is successful):

No feedback

If the command is not successful, this error message occurs:

CMMVC5753E The specified object does not exist or is not a suitable candidate.

setlocale

Use the **setlocale** command to change the locale setting for the clustered system (system). It also changes command output to the chosen language.

Syntax

```
►► — setlocale — — -locale — locale_id —————►◄
```

Parameters

-locale *locale_id*

Specifies the locale ID. The value must be a numeric value depending on the desired language (as indicated below)

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 system; all interface output is changed to the chosen language. For example, to change the language to Japanese, type the following:

```
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
- 8 Korean
- 9 Portuguese (Brazilian)
- 10 Russian

Note: This command does not change the front panel display panel settings.

An invocation example (where 3 is Japanese)

```
setlocale -locale 3
```

The resulting output:

No feedback

An invocation example (where 8 is Korean)

```
setlocale -locale 8
```

The resulting output:

No feedback

svqueryclock

Use the **svqueryclock** command to return the date, time, and current time-zone of the clustered system (system).

Syntax

```
▶▶—svqueryclock—▶▶
```

Parameters

None

Description

This command returns the date, time and current time-zone of the system.

An invocation example

```
svqueryclock
```

The resulting output:

```
Mon Nov 25 14:59:28 GMT 2013
```

writesernum

Use the **writesernum** command to write the node serial number into the planar NVRAM.

Syntax

```
▶▶  writesernum  —  —  -sernum  —  serial_number  —  [ node_id  ]
                                           [ node_name  ]▶▶
```

Parameters

-sernum *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 **lsnodevpd** command to verify that the number is correct. The `system_serial_number` field contains the serial number.

An invocation example

```
writesernum -sernum 1300027 node1
```

The resulting output:

No feedback


Chapter 9. Controller command

Use the controller command to modify the name of a storage controller.

chcontroller

Use the **chcontroller** command to modify the attributes of a controller.

Syntax



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.

-site site_name | site_id

(Optional) Specifies the numeric site value or site name for the controller. The value is 1, 2, or 3.

Note: The controller site cannot be changed if the system topology is stretched or hyperswap (and there are MDisk).

-nosite

(Optional) Resets the site value for the controller.

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 **lscontroller** command, the set quorum action fails for that MDisk. Before using the **chcontroller** command to set the **-allowquorum** parameter to **yes** on any disk controller, check the following website to see whether the controller supports quorum.

www.ibm.com/support

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 clustered system (system).

To add a new disk controller system to a running configuration, ensure that the system has detected the new storage MDisks by issuing the **detectmdisk** command. The controller has automatically been assigned a default name. If you are unsure of which controller is presenting the MDisks, issue the **lscontroller** 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.

Give this controller a descriptive name by issuing the following command:

```
chcontroller -name newname oldname
```

List the unmanaged MDisks by issuing the following command:

```
lsmdisk -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 storage pool and add only the RAID arrays that belong to the new controller to this storage pool. Avoid mixing RAID types; for each set of RAID array types (for example, RAID-1 or RAID-10), create a new storage pool. (You cannot use RAID-10 with distributed arrays.) Assign this storage pool an appropriate name; if your controller is called FAST650-abc and the storage pool contains RAID-10 arrays, assign the MDisk a name similar to F600-abc-R5. Issue the following command:

```
mkmdiskgrp -ext 16 -name mdisk_grp_name  
-mdisk colon-separated list of RAID-x mdisks returned
```

Note: This creates a new storage pool with an extent size of 16 MB.

An invocation example

```
chcontroller -name newtwo 2
```

The resulting output:

No feedback

An invocation example

```
chcontroller -site site1 controller18
```

The resulting output:


No feedback

lscontroller

Use the **lscontroller** command to display a concise list or a detailed view of controllers that are visible to the clustered system (system).

Syntax

```
➤— lscontroller — [ -filtervalue — attribute=value ] [ -nohdr ] —————➤
```



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.

-filtervalue?

(Optional) Displays the valid filter attributes. The following filter attributes for the **lscontroller** command are valid:

- controller_name
- id
- site_id
- site_name

-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.

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 system.

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 system, list the storage controllers by issuing the **lscontroller** command. Record the controller name or ID for the controller that you want to determine. For the controller in question, issue the **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 MDisk currently associated with this storage 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.

Remember: This value is reset by specific maintenance procedures or when the event log is cleared.

3. A SAN connection from a node or node canister port to a controller port for a single MDisk is a path. The controller port *path_count* value is the number of paths that are currently being used to submit input/output (I/O) data to this controller port.
4. The storage controller *max_path_count* value is the highest value that the storage controller *path_count* has reached since it was last reset to the *path_count* value. This value is reset by specific maintenance procedures or when the system error log is cleared.

Important: The *max_path_count* value is the highest value that the *path_count* has reached since it was last reset to the *path_count* value.

Remember: This value is reset by specific maintenance procedures or when the event log is cleared.

5. The *allow_quorum* value identifies if 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.

Important: This data comes from vendor-controlled sources and might not be available.

Table 59 provides the attribute values that can be displayed as output view data.

This table provides the attribute values that can be displayed as output view data.

Table 59. Iscontroller output

Attribute	Possible Values
id	Indicates the controller ID.
name	Indicates the controller name.
WWNN	Indicates the worldwide node name (WWNN). This field is blank for iSCSI controllers.
mdisk_link_count	Indicates the MDisk link count.
max_mdisk_link_count	Indicates the maximum MDisk link count.
degraded	Indicates if the controller has degraded MDisk.
vendor_id	Indicates the vendor identification name or number.
product_id_low	Indicates the product identification.
product_id_high	Indicates the product identification.
product_revision	Indicates the product revision.
ctrl_s/n	Indicates the controller serial number.
allow_quorum	Indicates the controller can support quorum disks.
WWPN	Indicates the worldwide port name (WWPN). This field is blank for iSCSI controllers.

Table 59. *lscontroller* output (continued)

Attribute	Possible Values
path_count	Indicates the number of paths that are currently being used to submit input/output (I/O) data to the controller port.
max_path_count	Indicates the maximum number of paths that are currently being used to submit input/output (I/O) data to the controller port.
site_id	Indicates the site value for the controller. This numeric value is 1, 2, 3 or blank.
site_name	Indicates the site name for the controller. This is an alphanumeric value or is blank.
fabric_type	Indicates a Fibre Channel (FC) or SAS controller. <ul style="list-style-type: none"> fc indicates an FC controller sas_direct indicates an SAS direct-attached controller multiple indicates multiple controllers (either FC, SAS, or both) iscsi indicates an iSCSI controller
iscsi_port_id	Indicates the I/O port identifier, which is the same as the WWPN value from the FC domain. This value shows the iSCSI port ID for the iSCSI controller and is blank for other controllers. This value must be a numeric value. This ID refers to the row number in output from the lsiscsistorageport command. The lsiscsistorageport can be used to find the controller IQN.
ip	Indicates the IP address associated with the <code>iscsi_port_id</code> . This shows the IP value for the iSCSI controller and is blank for other controllers. This must be an IPv4 or IPv6 address.

A concise invocation example for an iSCSI controller

```
lscontroller -delim :
```

The concise resulting output:

```
id:controller_name:ctrl_s/n:vendor_id:product_id_low:product_id_high:WWNN:degraded:fabric_type:site_id:site_name
0:controller0::IBM:1726-4xx:FAStT::no:iscsi:1:snpp1
1:controller1::IBM:1726-4xx:FAStT::no:iscsi:2:snpp2
7:controller7:3EK0J5Y8:SEAGATE :ST373405:FC:200600A0B851061E:yes:fc:1:snpp1
8:controller8:3EK0J6CR:SEAGATE :ST373405:FC:200600A0B851061D:no:fc:2:snpp2
9:controller9:3EK0J4YN:SEAGATE :ST373405:FC:200600A0B851061C:no:fc:3:snpp3
10:controller10:3EK0GKGH:SEAGATE :ST373405:FC:200600A0B851061B:no:fc:1:snpp4
11:controller11:3EK0J85C:SEAGATE :ST373405:FC:200600A0B851061A:no:fc:2:snpp5
12:controller12:3EK0JBR2:SEAGATE :ST373405:FC:200600A0B851062A:no:fc:3:snpp6
13:controller13:3EKYNJF8:SEAGATE :ST373405:FC:200600A0B851062B:no:fc:1:snpp7
14:controller14:3EK0HVTM:SEAGATE :ST373405:FC:200600A0B851062C:no:fc:2:snpp8
```

A detailed invocation example for a Fibre Channel controller

```
lscontroller -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
site_id=2
site_name=DR
WWPN=22000004CF2412AC
path_count=1
max_path_count=1
WWPN=21000004CF2412AC
path_count=0
max_path_count=0
fabric_type=sas_direct
iscsi_port_id=
ip=

```

A detailed invocation example for an iSCSI controller

lscontroller 0

The detailed resulting output:

```

id 0
controller_name controller0
WWNN
mdisk_link_count 4
max_mdisk_link_count 4
degraded no
vendor_id IBM
product_id_low 1726-4xx
product_id_high FAST
product_revision 0617
ctrl_s/n
allow_quorum no
fabric_type iscsi
site_id
site_name
WWPN
path_count 4
max_path_count 4
iscsi_port_id 1
ip 10.10.10.1
WWPN
path_count 4
max_path_count 4
iscsi_port_id 2
ip 10.10.10.2

```

lscontrollerdependentvdisks


Use the **lscontrollerdependentvdisks** command to list the volumes that are dependent on the specified controller.

Syntax

```

➔➔ lscontrollerdependentvdisks — [ -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.

<i>controller id list</i>	<i>controller name list</i>
---------------------------	-----------------------------

Specifies one or more controller IDs, controller names, or both. Separate multiple controllers using the colon character (:).

Description

The **lscontrollerdependentvdisks** command lists the volumes that are dependent on the status of the specified controllers. If a controller goes offline, the dependent volumes 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 volumes.

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 volumes.

Note: The command lists the volumes 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

```
lscontrollerdependentvdisks controller0
```

The concise resulting output:

```
vdisk_id vdisk_name
0 vdisk0
1 vdisk1
2 vdisk2
```

Chapter 10. Drive commands

Use the drive commands to capture information to assist with managing drives.

applydrivesoftware

Use the **applydrivesoftware** command to update drives.

Syntax

```
►► applydrivesoftware — -file name [ -type fpga | firmware ] [ -drive drive_id | -all ]
[ -force ] [ -allowreinstall ] [ -allowdowngrade ]
►► applydrivesoftware — -cancel
```

Parameters

-file *name*

(Required) Specifies the firmware update file name that must be copied to the /home/admin/update/ directory on the configuration node.

-type *fpga* | *firmware*

(Optional) Specifies the type of drive firmware to update. Drive firmware updates can be performed online, concurrently with I/O. However, fpga updates require the drive to be taken offline, which means target drives must be made candidate before issuing the **applydrivesoftware** command. The default value is *firmware*. See the **chdrive** command for more details.

-all

(Optional) Specifies that the drive firmware should be applied to every drive in the system, as long as that drive is online and has use member, use spare, or use candidate.

This does not apply to:

- Drives that have dependent volumes
- Drives that are members of non-redundant arrays

Drives hosting quorum qualify, but there is risk. To avoid this risk use **-drive** and make sure the quorum is moved in between **applydrivesoftware** invocations. Use the **chquorum** command to avoid updating a drive that is hosting quorum.

If you specify **-all** you must specify the **-type** as *firmware*.

Remember: The **-all** parameter differs from the **-drive** parameter because unsuitable drives are not added to the list of drives scheduled for update when you use **-all**.

-drive *drive_id*

(Optional) Specifies one drive ID or a list of drive IDs (separated by a colon, [:]) to be updated. The maximum number of IDs is 128. If you have more than 128, use **-all** or multiple **applydrivesoftware** invocations to complete the update.

Remember: The **-drive** parameter differs from the **-all** parameter because if you specify an unsuitable drive using the **-drive** parameter, the **applydrivesoftware** command fails. Additionally, if you specify all three drives as hosting quorum, the command fails. If you use the **-drive** option to specify a single drive, and that drive has `use=unused`, it is updated.

-force

(Optional) Bypasses the dependent volume check. By default **applydrivesoftware** cannot run if any volumes dependent on the drive. Specifying **-force** bypasses this check, allowing the drive software update to proceed. Drive-dependent volumes generally result from non-redundant or degraded RAID arrays.

Note: Restore redundancy to the system (where possible) instead of using the **-force** parameter.

Important: Using the **-force** parameter might result in a data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of data loss in the array or pool to which the drive belongs.

-allowreinstall

(Optional) Specifies to make the system install the current level (again) onto drives that contain a file in the package.

Remember: Using this parameter is not recommended.

-allowdowngrade

(Optional) Specifies to allow the system to downgrade the firmware on a drive (that contains a file in the package).

Remember: Use this parameter only under the direction of your product support information.

-cancel

(Optional) Specifies that the command be stopped.

Description

Use this command to update the firmware of drives that are managed by the system.

There are two types of drive software that can be updated using this command:

- firmware
- fpga

Drive firmware updates can be performed online while the drive is in use. When used on an array member drive **applydrivesoftware** checks for volumes that are dependent on the drive and refuses to run if any are found. Drive dependent volumes are usually caused by non-redundant or degraded RAID arrays. Where possible you should restore redundancy to the system by replacing any failed drives before using the **applydrivesoftware** command. When this is not possible, for example on drives that are members of a RAID-0 array, you can either add redundancy to the volume by adding a second copy in another pool, or use the **-force** parameter to bypass the dependant volume check.

Remember: Only use **-force** if you are willing to accept the risk of data loss on dependent volumes (if the drive fails during the firmware update).

Drive firmware updates occur asynchronously, and conclude after the **applydrivesoftware** command completes. To see the status of the updates, use the **lsdriveupgradeprogress** command.

Drive fpga updates might require the drive to be taken offline for several minutes. Drives must be changed to the candidate state before **applydrivesoftware** can be used to update fpga software. The fpga

updates occur asynchronously, continuing in the background after the **applydrivesoftware** command has returned. You must check the `FPGA_level` field in **lsdrive** *N*, where *N* is the *drive_id*, to see whether or not the update completed successfully.

Remember: Interrupting an fpga update by removing power from the drive or enclosure might make the drive unusable. Only one drive's fpga can be updated per **applydrivesoftware** invocation. Make sure that the update is complete before unseating the drive or removing power from the enclosure.

An invocation example

```
applydrivesoftware -file DRIVE_XXXXXXX -type firmware -drive 4
```

The resulting output:

No feedback

An invocation example

An example that fails because of drive-dependent volumes:

```
applydrivesoftware -file DRIVE_XXXXXXX -type firmware -drive 6
```

The resulting output:

```
CMMVC6953E The action cannot be completed because vdisks are dependent on the specified mdisk.  
Force is required.
```

An invocation example

```
applydrivesoftware -file drivemicrocodepackagev5 -type firmware -all
```

The resulting output:

No feedback

An invocation example

```
applydrivesoftware -file drivemicrocodepackagev1 -type firmware -all -allowreinstall
```

The resulting output:

No feedback

An invocation example

```
applydrivesoftware -file drivemicrocodepackagev1 -type firmware -all -allowdowngrade
```

The resulting output:

No feedback

An invocation example

```
applydrivesoftware -file drivemicrocodepackagev1 -type firmware -all -allowdowngrade  
-allowreinstall
```

The resulting output:

No feedback

An invocation example

```
applydrivesoftware -cancel
```

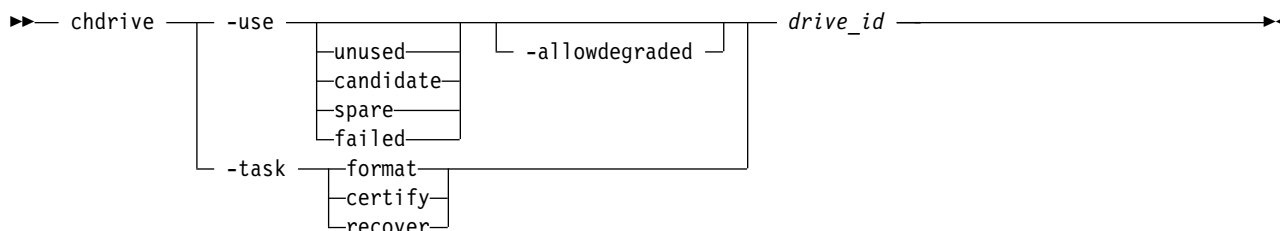
The resulting output:

No feedback

chdrive

Use the **chdrive** command to change the drive properties.

Syntax



Parameters

-use unused | candidate | spare | failed

Describes the role of the drive:

- **unused** indicates the drive is not in use and will not be used as a spare
- **candidate** indicates the drive is available for use in an array
- **spare** indicates the drive can be used as a hot spare if required
- **failed** indicates the drive has failed.

Note: To create member drives, add the drives to (new) arrays using the **mkarray** command.

If a drive fails for a distributed array, the array remains associated with the failed drive while it is in the failed state.

-allowdegraded

(Optional) Permits permission for a change of drive use to continue, even if a hotspare drive is not available for the array that the drive is a member of. You cannot specify **-allowdegraded** and **-task** together.

Important: Using **-allowdegraded** is not recommended.

-task format | certify | recover

Causes the drive to perform a task:

- **format** indicates a drive will be formatted for use in an array; only permitted when drive is a candidate or has failed validation.
- **certify** indicates the drive will be analyzed to verify the integrity of the data it contains; permitted for any drive that is a member.
- **recover** recovers an offline flash drive without losing data; permitted when the drive is offline because a build is required, or when the drive has failed validation.

drive_id

The identity of the drive.

Description

Use this command to change the drive role, or to start long running drive tasks.

You can use **lsdriveprogress** to display progress (percentage) and estimated completion time of ongoing drive tasks.

When a drive associated with a distributed array is changed from member to failed, if the distributed array does not have available rebuild space then it is degraded. If **-allowdegraded** is not specified the command fails because of insufficient rebuild areas. If the **-allowdegraded** parameter is specified the command succeeds and the array no longer uses the drive for I/O operations. If a drive is changed from failed to another configuration the distributed array forgets about the drive and creates a missing member that belongs in the member table. Use the **chararraymember** command to replace the missing member.

An invocation example

```
chdrive -use spare 1
```

The resulting output:

No feedback

An invocation example to certify drive 23

```
chdrive -task certify 23
```

The resulting output:

No feedback

lsdrive

Use the **lsdrive** command to display configuration information and drive vital product data (VPD).

Syntax

```
►► lsdrive — [ -filtervalue — attribute_value ] [ -nohdr ] [ -delim — delimiter ] [ -filtervalue? ] [ -bytes ] [ drive_id ] —►
```

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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""): **lsdrive -filtervalue mdisk_id="1*"**

-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?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- capacity
- enclosure_id
- error_sequence_number
- id
- interface_speed
- mdisk_id
- mdisk_name
- member_id
- node_id
- node_name
- slot_id
- status
- tech_type
- use
- drive_class_id

-bytes

(Optional) The size (capacity) of the drive in bytes.

drive_id

(Optional) 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 60 on page 345 describes possible outputs.

Table 60. Isdrive output

Attribute	Value
id	Indicates the ID of the drive: <ul style="list-style-type: none"> • online, which indicates that the drive is available through all drive ports • degraded, which indicates that the drive is available but not through all drive ports • offline, which indicates that the drive is unavailable
status	Indicates the summary status of the drive.
error_sequence_number	Indicates the error sequence number describing the cause of the drive status: <ul style="list-style-type: none"> • online, which is blank • degraded, which is populated if associated with an error • offline, which 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	Indicates the current role of the drive: <ul style="list-style-type: none"> • unused if the drive which is not configured to be used by anything • candidate if the drive is available to be configured • spare if the drive is configured as a spare, to be used if the arrays fail members • member if the drive is configured as a member of an array • failed if the drive has been rejected and is no longer available for use
UID	Indicates the unique ID reported by the drive.
tech_type	Indicates the drive technology used. <p>The values are:</p> <ul style="list-style-type: none"> • unsupported indicates that the drive is not supported in this platform (contact your support team). • tier0_flash • tier1_flash • tier_enterprise • tier_nearline • drive.flashcard indicates the flash drive.
replacement_date	Indicates the date of a potential drive failure. The format must be YYMMDD. This is blank for non-SSD drives.
capacity	Indicates the capacity of disk, excluding quorum area.
block_size	Indicates the block size of the disk.
vendor_id	Indicates the manufacturer of the drive.
product_id	Indicates the product ID of the drive.
FRU_part_number	Indicates the FRU part number of the drive.
FRU_identity	Indicates the 11S number combining manufacturing part number and serial number.
RPM	Indicates the specified RPM of the disk.
firmware_level	Indicates the firmware level of the disk; blank if unknown.
FPGA_level	Indicates the FPGA level, if applicable; blank if not applicable or unknown.
mdisk_id	Indicates the ID of the array MDisk that the drive is a member of.
mdisk_name	Indicates the name of the MDisk that the drive is a member of.
member_id	Indicates the ID of the MDisk array member.

Table 60. Isdrive output (continued)

Attribute	Value
enclosure_id	Indicates if the: <ul style="list-style-type: none"> • 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. • Drive is contained in a node (not an enclosure), this is blank. • Enclosure ID has not been determined yet, this is blank.
slot_id	Indicates 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	Indicates the node name where the drive is located for a drive contained within a node. For a drive contained within an enclosure, it is blank.
node_id	Indicates the node ID where the drive is located for a drive contained within a node. For a drive contained within an enclosure, blank.
quorum_id	Indicates the ID of quorum disk; blank if not quorum disk.
port_1_status	Indicates the connectivity status of the drive's first port. The values are online, offline, or excluded. Note: Port 1 is attached to the node that has a panel name ending in -1.
port_2_status	Indicates the connectivity status of the drive's second port. The values are online, offline, or excluded. Note: Port 2 is attached to the node that has a panel name ending in -2.
interface_speed	Indicates the lowest interface speed for the connected drive slot (in gigabits per second, or Gbps). The values are: <ul style="list-style-type: none"> • 1.5 Gbps • 3 Gbps • 6 Gbps • 12 Gbps • Blank if both ports are isolated or the drive is not connected
protection_enabled	Indicates if SCSI type-2 protection information is enabled (yes) or not (no).
auto_manage	Indicates if the auto_manage process is running (active) or not running (idle).
drive_class_id	Indicates which drive class the drive is part of.
write_endurance_used	Indicates the drive writes per day (DWPD). This value is blank for drives that are not SSD drives. The value must be a number 0 - 255. This value indicates the percentage of life used by the drive. The value 0 indicates that full life remains, and 100 indicates that the drive is at or past its end of life. Note: The drive must be replaced when the value exceeds 100. This value is blank for drives that are either: <ol style="list-style-type: none"> 1. Not SSDs. 2. SSDs that predate support of the endurance indicator. This also applies to drives that are yet to be polled, which can take up to 24 hours.
write_endurance_usage_rate	Indicates the DWPD usage rate. The values are: <ul style="list-style-type: none"> • measuring • high • marginal • low This is blank for non-SSD drives.

A concise invocation example

```
lsdrive -delim :
```

The concise resulting output:

```
id:status:error_sequence_number:use:tech_type:capacity:mdisk_id:mdisk_name:member_id:enclosure_id:slot_id:auto_manage:dr
0:online::member:tier0_flash:20GB:0:mdisk0:0:1:2:active:0
1:offline:345:member:tier0_flash:20GB:0:mdisk0:0:1:3:idle:0
2:online::member:tier0_flash:20GB:0:mdisk0:0:1:4:active:0
```

A detailed invocation example for an SSD drive

```
lsdrive 0
```

The detailed resulting output:

```
id:0
status:online
error_sequence_number:
use:member
UID:20000004cf4cd2c0

tech_type:tier0_flash
capacity:20GB
block_size:512
vendor_id:IBM
product_id:I8MR1337 W00Y4Y1
FRU_part_number:AAAAAAA
FRU_identity:11S1817115Y41337171001
RPM:
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
interface_speed:6Gb
protection_enabled:yes
auto_manage:active
drive_class_id:3
write_endurance_used:5
write_endurance_usage_rate:high
work_load:high
replacement_date:190806
```

A detailed invocation example for a tier 1 flash SSD drive

```
lsdrive 0
```

The detailed resulting output:

```
id 0
status degraded
error_sequence_number
use candidate
UID 5000c5002624a723
tech_type sas_hdd
capacity 1.8TB
block_size 512
vendor_id IBM-207x
product_id ST32000444SS
FRU_part_number 85Y5869
```

```

FRU_identity 11S41Y8471YXXX9WM40LMD
RPM 10000
firmware_level BC2D
FPGA_level
mdisk_id
mdisk_name
member_id
enclosure_id 1
slot_id 7
node_id
node_name
quorum_id 0
port_1_status online
port_2_status offline
interface_speed 6Gb
protection_enabled no
auto_manage inactive
drive_class_id 3
write_endurance_used 30
drive_class_id
write_endurance_used 5
write_endurance_usage_rate high
work_load high
replacement_date 190806

```

A detailed invocation example for a hard disk drive (HDD)

```
lsdrive 0
```

The detailed resulting output:

```


id 0
status degraded
error_sequence_number
use_candidate
UID 5000c5002624a723
tech_type sas_nearline_hdd
capacity 1.8TB
block_size 512
vendor_id IBM-207x
product_id ST32000444SS
FRU_part_number 85Y5869
FRU_identity 11S41Y8471YXXX9WM40LMD
RPM 7200
firmware_level BC2D
FPGA_level
mdisk_id
mdisk_name
member_id
enclosure_id 1
slot_id 7
node_id
node_name
quorum_id 0
port_1_status online
port_2_status offline
interface_speed 6Gb
protection_enabled no
auto_manage inactive
drive_class_id 3
write_endurance_used
drive_class_id
write_endurance_used 5
write_endurance_usage_rate high
work_load high
replacement_date 190806

```

lsdriveclass

Use the **lsdriveclass** command to display all drive classes in the clustered system (system).

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.

-filtervalue *attribute=value*

(Optional) Specifies a list of one or more filter attributes that match the specified values; see **-filtervalue?** for the supported attributes.

Note: Some filters allow the use of a wildcard when entering the command. The following rules apply to the use of wildcards when using the 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:

```
lsdriveclass -filtervalue "IO_group_name=md*"
```

-filtervalue?

(Optional) Includes all of the valid filter attributes in the report. The following filter attributes are valid for the **lsdriveclass** command:

- id
- RPM
- capacity
- IO_group_id
- IO_group_name
- tech_type
- block_size
- candidate_count
- superior_count
- total_count

Any parameters that are specified with the **-filtervalue?** parameter are ignored.

-delim *delimiter*

(Optional) By default, in a concise view all columns of data are space-separated, with the width of each column set to the maximum width of each item of data. In a detailed view, each item of data is an individual row, and if you display headers, the data is separated from the header by a space. The

-delim parameter overrides this behavior. Valid input for the **-delim** parameter is a 1-byte character. Enter **-delim :** on the command line, and 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 specified *delimiter* separates the data from its header.

drive_class_id

(Optional) The identity of the drive class.

Description

This command displays all drive classes in a system. Drives are displayed if they are managed.

Table 61 provides the attribute values that can be displayed as output view data.

Table 61. lsdriveclass output

Attribute	Possible Values
id	Indicates the drive class ID.
RPM	Indicates the speed of the drive class.
capacity	Indicates the capacity of the drive class.
IO_group_id	Indicates the I/O group ID associated with the drive class.
IO_group_name	Indicates the I/O group name that is associated with the drive class.
tech_type	Indicates the technology type of the drive class.
block_size	Indicates the block size of the drive class.
candidate_count	Indicates the number of drives in the drive class that are in candidate state.
superior_count	Indicates the total number of drives in this class and those that count as superior. (This applies to distributed arrays created using mkdistributedarray .)
total_count	Indicates the total number of drives in the drive class. The drive state is irrelevant.

A concise invocation example

```
lsdriveclass -filtervalue block_size=4096
```

The detailed resulting output:

id	RPM	capacity	IO_group_id	IO_group_name	tech_type	block_size	candidate_count	superior_count	total_count
3	15000	600.5GB	2	io_group2	tier0_flash	4096	0	0	24

A concise invocation example

```
lsdriveclass -filtervalue io_group_ID=0:tech_type=tier_enterprise
```

The detailed resulting output:

id	RPM	capacity	IO_group_id	IO_group_name	tech_type	block_size	candidate_count	superior_count	total_count
0	10000	300.9GB	0	io_group0	tier0_flash	512	30	3	30

A concise invocation example

```
lsdriveclass -delim ! -nohdr
```

The detailed resulting output:


```
0!10000!300.9GB!0!io_group0!tier_nearline!512!30!30
1!1600.5GB!0!io_group0!tier_nearline!512!10!50
2!15000!900.1GB!1!io_group1!tier_enterprise!512!60!60
3!15000!600.5GB!2!io_group2!tier_enterprise!4096!0!24
```

A detailed invocation example

```
lsdriveclass 2
```

The detailed resulting output:

```
id 2
RPM 15000
capacity 900.1GB
IO_group_id 1
IO_group_name io_group1
tech_type tier0_flash
block_size 512
candidate_count 60
superior_count 5
total_count 60
```

lsdrivelba

Use the **lsdrivelba** command to map array MDisk logical block address (LBA) to a set of drives.

Syntax

```

>>> lsdrivelba [-nohdr] [-delim delimiter] -mdisklba lba
>>> -mdisk 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.

-mdisklba lba

(Required) The logical block address (LBA) on the MDisk. The LBA must be specified in hex, with a 0x prefix.

-mdiskmdisk_id | mdisk_name

(Required) The ID or name of the MDisk.

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.

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 (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:
`lsdriveprogress -filtervalue "task=*`

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalueattribute=value** parameter:

- task

drive_id

(Optional) The drive for which you want to view progress.

Description

The following outputs are possible:

drive_id

Indicates the ID for the drive with the active task.

task Indicates the type of task:

- format
- certify
- recover

progress

Indicates the percentage complete of the job.

estimated_completion_time

Indicates the estimated completion time, in the format YYMMDDHHMMSS, where:

- Y is year
- (The first) M is month
- D is day
- H is hour
- (The second) M is minute
- S is second

.

An invocation example

```
lsdriveprogress -delim :
```

The resulting output:

```
drive_id:task:progress:estimated_completion_time
0:format:10:091118131056
9:certify:25:991231235959
```

An invocation example

```
lsdriveprogress -delim : 9
```

The resulting output:

```
9:certify:25:991231235959
```

lsdriveupgradeprogress

Use the **lsdriveupgradeprogress** command to view the status or progress of drives with pending downloads.

Syntax

```
➤—lsdriveupgradeprogress—┬── -delim — delimiter ─┬── drive_id ─┬──➤
```

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.

drive_id

(Optional) Specifies the update status or progress for a single drive. If not specified, the update status for all scheduled drives is displayed.

Note: If you specify this parameter, **lsdriveupgradeprogress** displays the update status of this drive. If you do not specify this parameter, **lsdriveupgradeprogress** displays the update status of all requested drives.

Description

The **lsdriveupgradeprogress** command completes whether the original **applydrivesoftware** command was addressed to a single drive (synchronous command) or multiple drives (asynchronous command).

The following outputs are possible:

id Indicates the identity of the active drive.

status Indicates the drive status. Each of the following values has a specific meaning:

- progressing indicates all scheduled drives have completed; wait 270 seconds before you issue **applydrivesoftware** again.
- completed indicates a successful firmware download.
- updating indicates the update is ongoing.
- scheduled indicates the update is in the download list, waiting to download.
- not_scheduled indicates the drive is not scheduled.

Note: This means the corresponding drive was not scheduled when **applydrivesoftware -all** was last issued.

- canceled indicates the update has been canceled, either by command or by a system change that stops the **applydrivesoftware** command from running properly.
- invalid indicates the drive download status is invalid, also the initial state.
- If the field is blank, that indicates that the download is not scheduled (for example, the **applydrivesoftware** command has not been issued).

estimated_completion_time

Indicates the estimated completion time (YYMMDDHHMMSS), where:

- Y is year
- (The first) M is month
- D is day
- H is hour
- (The second) M is minute
- S is second

The value is blank if the status is either canceled or blank.

A concise invocation example with two drives canceled

```
lsdriveupgradeprogess
```

The resulting output:

```
id status    estimated_completion_time
0  completed  121112062608
5  canceled
6  canceled
```

A concise invocation example with drives scheduled or completed

```
lsdriveupgradeprogess
```

The resulting output:

```
id status    estimated_completion_time
0  completed  121112062608
5  scheduled  121112062638
6  scheduled  121112062708
```

A detailed invocation example using a drive ID

```
lsdriveupgradeprogess 17
```

The resulting output:

```
id status    estimated_completion_time
17 completed  121123134627
```

A concise invocation example

```
lsdriveupgradeprogess -delim :
```

The resulting output:

```
id:status:estimated_completion_time
0:completed:121101065019
1:scheduled:121101065049
2:scheduled:121101065119
```

A concise invocation example

```
lsdriveupgradeprogess
```

The resulting output:

```
id status      estimated_completion_time
24 completed 121212164752
25 canceled
26 canceled
```

A concise invocation example

`lsdriveupgradeprogress`

The resulting output:

```
id status      estimated_completion_time
0  completed 130714223913
1  completed 130714223943
2  completed 130714224013
3  completed 130714224043
4  completed 130714224113
5  completed 130714224143
6  completed 130714224213
7  completed 130714224243
8  completed 130714224313
9  completed 130714224343
10 completed 130714224413
11 completed 130714224443
```

A concise invocation example

`lsdriveupgradeprogress -delim :`

The resulting output:

```
id:status:estimated_completion_time
0:completed:130714223913
1:completed:130714223943
2:completed:130714224013
3:completed:130714224043
4:completed:130714224113
5:completed:130714224143
6:completed:130714224213
7:completed:130714224243
8:completed:130714224313
9:completed:130714224343
10:completed:130714224413
11:completed:130714224443
```

triggerdrivedump

Use the **triggerdrivedump** command to collect support data from a disk drive. This data can help you understand problems with the drive, and does not contain any data that applications might have written to the drive.

Syntax

►► — `triggerdrivedump —drive_id` —◀◀

Parameters

drive_id

(Required) The ID of the drive to dump.

Description

Use this command to collect internal log data from a drive and store the information in a file in the /dumps/drive directory. This directory is on one of the nodes connected to the drive. The system limits the number of drive dump files in the directory to 24 per node.

An invocation example

```
triggerdrivedump 2
```

The resulting output:

```
Dump file for drive [2] created
```

Note: The system chooses the node on which to run the statesave.

Chapter 11. Email and event notification commands

You can use the command-line interface (CLI) to enable your system to send notifications.

chemail

Use the **chemail** command to set or modify contact information for email event notifications. At least one of the parameters must be specified to modify settings.

Syntax

```
➤ 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 ] [ -organization — organization ] [ -address — address ]
➤ [ -city — city ] [ -state — state ] [ -zip — zip ] [ -country — country ]
```

Parameters

-reply *reply_email_address*

(Optional) Specifies the email address to which a reply is sent.

-contact *contact_name*

(Optional) Specifies the name of the person to receive the email.

For machine types 2071 and 2072 the maximum number of characters is 30. For other machine types the maximum number of characters is 72.

-primary *primary_telephone_number*

(Optional) Specifies the primary contact telephone number.

Note: For machine types 2071 and 2072 (in the United States and Canada), the value entered must be exactly ten decimal digits. For machines types 2071 and 2072 (in other countries) the value entered can be five to nineteen decimal digits. Otherwise, there can be up to nineteen characters.

-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 email.

For machine types 2071 and 2072 the maximum number of characters is 30. For other machine types the maximum number of characters is 72.

-primary2 *primary_telephone_number2*

(Optional) Specifies the primary contact telephone number for the second contact person.

Note: For machine types 2071 and 2072 (in the United States and Canada), the value entered must be exactly ten decimal digits. For machines types 2071 and 2072 (in other countries) the value entered can be five to nineteen decimal digits. Otherwise, there can be up to nineteen characters.

-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.

-organization *organization*

(Optional) Specifies the user's organization as it should appear in Call Home emails.

-address *address*

(Optional) Specifies the first line of the user's address as it should appear in Call Home email.

-city *city*

(Optional) Specifies the user's city as it should appear in Call Home email.

-state *state*

(Optional) Specifies the user's state as it should appear in Call Home email. This is a two-character value such as NY for New York.

-zip *zip*

(Optional) Specifies the user's zip code or postal code as it should appear in Call Home email.

-country *country*

(Optional) Specifies the country in which the machine resides as it should appear in Call Home email. This is a two-character value such as US for United States.

For machine types 2071 and 2072 this value cannot be US or CA if the value for **primary** or **primary2** telephone number is not blank or exactly 10 digits.

Description

This command sets or modifies contact information that is used by the email event notification facility.

Note: If you are starting the email event notification facility, the **reply**, **contact**, **primary**, and **location** parameters are required. If you are modifying contact information used by the email event notification facility, at least one of the parameters must be specified.

Remember: When considering e-mail addresses:

- Alphanumeric characters and additionally underscore (_), at symbol (@), and dot (.) characters are permitted.
- There must be exactly one @ character in the string, and the @ characters must not start or end the string.
- A plus (+) character is permitted before the @ character.

These fields do not have to be set to start the email notification system, but if the new fields are set they are included in the email event notifications.

An invocation example

```
chemail -reply ddrogba@uk.uefa.com
        -contact 'Didier Drogba'
        -primary 01962817668
        -location 'C block'
        -organization UEFA
        -address '1 Chelsea Blvd'
        -city Fulham
        -zip 0U812
        -machine_country GB
```

The resulting output:

No feedback

An invocation example

```
chemail -primary 0441234567 -location 'room 256 floor 1'
```

The resulting output:

No feedback

An invocation example

```
chemail -country US -primary 8458765309
```

The resulting output:

No feedback

chemailserver

Use the **chemailserver** command to modify the parameters of an existing email server object.

Syntax

```
➤➤ 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 email server object. The name must be a 1-through 63-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 email 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 email 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 email server object. The email server object describes a remote Simple Mail Transfer Protocol (SMTP) email server.

You must specify either the current name or the ID of the object returned at creation time. Use the **lsemailserver** command to obtain this ID.

An invocation example

```
chemailserver -name newserver 0
```

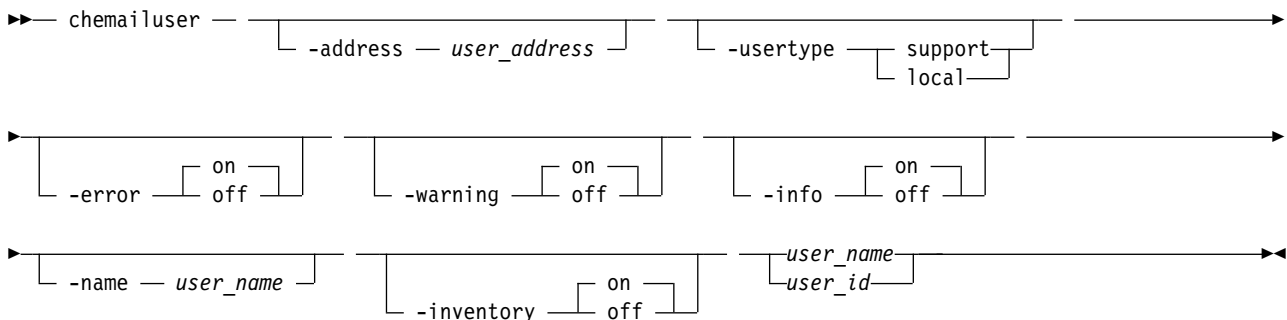
The resulting output:

No feedback

chemailuser

Use the **chemailuser** command to modify the settings that are defined for an email recipient.

Syntax



Parameters

-address *user_address*

(Optional) Specifies the email address of the person receiving the email 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 email 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 email 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 email recipient. Set to off, informational event notifications are not sent to the recipient.

-name *user_name*

(Optional) Specifies the user name of the new email 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 email notifications.

user_name* | *user_id

(Required) Specifies the email recipient for whom you are modifying settings.

Description

This command modifies the settings that are established for an email 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.

Remember: When considering e-mail addresses:

- Alphanumeric characters and additionally underscore (_), at symbol (@), and dot (.) characters are permitted.
- There must be exactly one @ character in the string, and the @ characters must not start or end the string.
- A plus (+) character is permitted before the @ character.

An invocation example

The following example modifies email settings for email recipient manager2008:

```
chemailuser -usertype local manager2008
```

The resulting output:

No feedback

An invocation example

The following example modifies email settings:

```
chemailuser -address fred@gmail.com -name Fred
```

The resulting output:


No feedback

chsnmpserver

Use the **chsnmpserver** command to modify the parameters of an existing SNMP server.

Syntax

```
►► chsnmpserver — [ -name — server_name ] [ -ip — ip_address ] —►
```



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 **lssnmpserver** command to obtain this ID.

An invocation example

```
chsnmpserver -name newserver 0
```


The resulting output:

```
No feedback
```

lsemailserver

Use the **lsemailserver** command to display a concise list or a detailed view of email servers configured on the clustered system (system).

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.

email_server_name | email_server_id

(Optional) Specifies the name or ID of an existing email server that must be listed.

Description

Use this command to display a concise list or a detailed view of email servers that are configured on the system.

A concise invocation example

```
lsmailserver -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

```
lsmailserver email0
```


The detailed resulting output:

```
id 0
name emailserver0
IP_address 192.135.60.3
port 25
```

lsemailer

Use the **lsemailer** command to generate a report that lists the email event notification settings for all email recipients, an individual email recipient, or a specified type (local or support) of email recipient.

Syntax



Parameters

-type support | local

(Optional) Specifies the types of email recipients you want to view, either customer-based 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.

user_name | user_id

(Optional) Specifies the user ID or user name of the email event recipient for whom you want to see the email notification settings.

Description

When you issue this command, a report is displayed that lists the email event notification settings for all email recipients, an individual email recipient, or a specified type (local or support) of email recipient. The concise and detailed views report the same information.

A concise invocation example listing information for all email recipients using the email event notification facility

```
lsemailuser -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
```

lssnmpserver

Use the **lssnmpserver** command to return a concise list or a detailed view of SNMP servers that are configured on the clustered system (system).

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 system.

A concise invocation example

```
lssnmpserver -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

```
lssnmpserver 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
```

mkemailserver

Use the **mkemailserver** command to create an email server object that describes a remote Simple Mail Transfer Protocol (SMTP) email server.

Syntax

```
➤ mkemailserver — -name — server_name — -ip — ip_address — -port — port ➤
```

Parameters

-name *server_name*

(Optional) Specifies a unique name to assign to the email server object. The name must be a 1-through 63-character string, and cannot start with a hyphen or number. If a name is not specified, then a system default of `emailserver n` 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 email 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 email server. This must be a value of 1 - 65535. The default value is 25.

Description

This command creates an email server object that represents the SMTP server. The SAN Volume Controller uses the email server to send event notification and inventory emails to email users. It can transmit any combination of error, warning, and informational notification types.

The SAN Volume Controller supports up to six email servers to provide redundant access to the external email network. The email servers are used in turn until the email is successfully sent from the SAN Volume Controller. The attempt is successful when the SAN Volume Controller gets a positive acknowledgement from an email server that the email has been received by the server.

An invocation example

```
mkemailserver -ip 2.2.2.2 -port 78
```

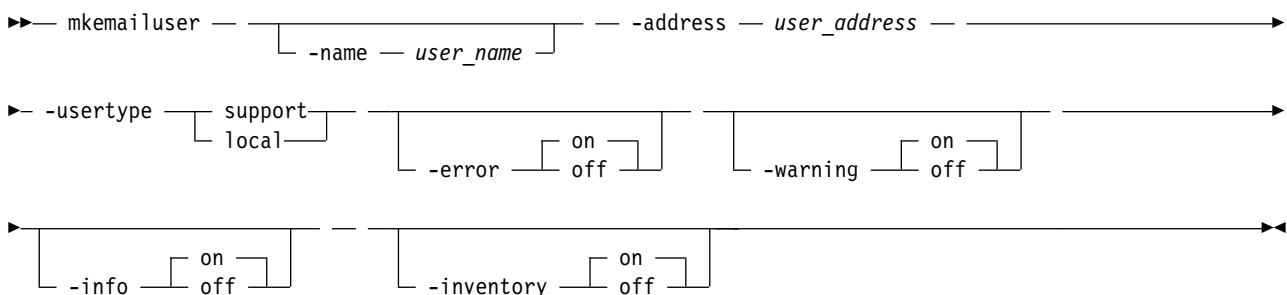
The resulting output:

```
Emailserver id [2] successfully created
```

mkemailuser

Use the **mkemailuser** command to add a recipient of email event and inventory notifications to the email event notification facility. Add up to twelve recipients (one recipient at a time).

Syntax



Parameters

-name *user_name*

(Optional) Specifies the name of the person who is the recipient of email 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 `emailuser n` , where n is a number beginning with 0 (`emailuser0`, `emailuser1`, and so on).

The name `emailuser n` , where n is a number, is reserved and cannot be used as one of your user names.

-address *user_address*

(Required) Specifies the email address of the person receiving the email event or inventory notifications, or both. The *user_address* value must be unique.

-usertype **support** | **local**

(Required) Specifies the type of user, either local or support, based on the following guidelines:

support

The recipient is your product support organization, using a default value (an auto-populated email address). This setting is used with the Call Home feature. For any other use case, contact your product support organization for direction.

local All other recipients other than your product support organization. Select the `local` usertype unless otherwise instructed by your product support organization.

-error on | **off**

(Optional) Specifies whether the recipient receives error-type event notifications. Set to `on`, error-type event notifications are sent to the email 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 email 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 email 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 email notifications. The default value is `off`.

Description

This command adds email recipients to the email event and inventory notification facility. You can add up to twelve recipients, one recipient at a time. When an email 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. Email notification starts when you process the **startemail** command.

Note: Before you can set the **usertype** parameter to `support`, turn the **-warning** and **-info** flags off.

Remember: When considering e-mail addresses:

- Alphanumeric characters and additionally underscore (`_`), at symbol (`@`), and dot (`.`) characters are permitted.
- There must be exactly one `@` character in the string, and the `@` characters must not start or end the string.
- A plus (`+`) character is permitted before the `@` character.

An invocation example

```
mkemailuser -address manager2008@ibm.com -error on -usertype local
```

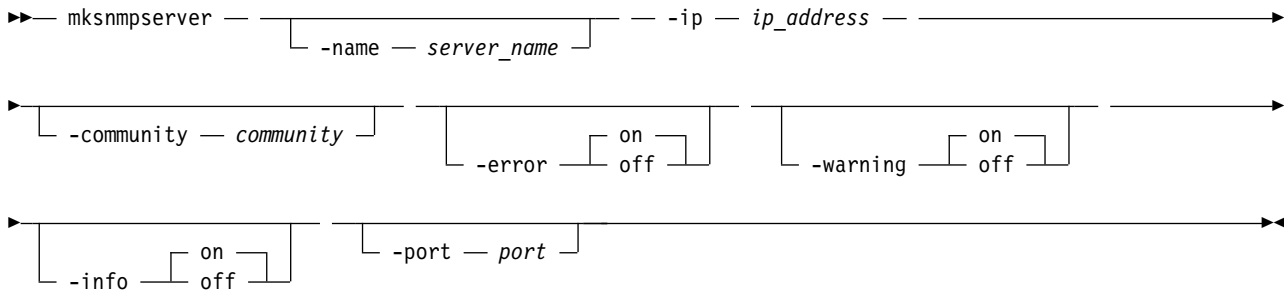
The resulting output:

email user, id [2], successfully created

mksnmpserver

Use the **mksnmpserver** command to create a Simple Network Management Protocol (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

```
mksnmpserver -ip 2.2.2.2 -port 78
```

The resulting output:

```
SNMP Server id [2] successfully created
```

rmemailserver

Use the **rmemailserver** command to delete the specified email server object.

Syntax

```
➤— rmemailserver — [ email_server_name | email_server_id ] —➤
```

Parameters

email_server_name | *email_server_id*

(Required) Specifies the name or ID of the email server object to be deleted.

Description

Use this command to delete an existing email server object that describes a remote Simple Mail Transfer Protocol (SMTP) email server. You must specify either the current name or the ID of the object returned at creation time. Use the **lserver** command to obtain this ID.

Note: Email service stops when the last email server is removed. Use the **startemail** command to reactivate the email and inventory notification function after at least one email server has been configured.

An invocation example

```
rmemailserver email4
```

The resulting output:

```
none
```

rmemailuser

Use the **rmemailuser** command to remove a previously defined email recipient from the system.

Syntax

```
➤— rmemailuser — [ user_name | user_id ] —➤
```

Parameters

user_name | *user_id*

(Required) Specifies the user ID or user name of the email recipient to remove.

Description

This command removes an existing email recipient from the system.

An invocation example to remove email recipient manager2008

```
rmemailuser manager2008
```

The resulting output:

No feedback

An invocation example to remove email recipient 2

```
rmemailuser 2
```

The resulting output:

No feedback

rmsnmpserver

Use the **rmsnmpserver** command to delete the specified Simple Network Management Protocol (SNMP)server.

Syntax

```
➤— 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 **lssnmpserver** command to obtain this ID.

An invocation example

```
rmsnmpserver snmp4
```

The resulting output:

No feedback

sendinventoryemail

Use the **sendinventoryemail** command to send an inventory email notification to all email recipients able to receive inventory email notifications. There are no parameters for this command.

Syntax

```
➤— sendinventoryemail —➤
```

Parameters

There are no parameters for this command.

Description

This command sends an inventory email notification to all email recipients who are enabled to receive inventory email notifications. This command fails if the **startemail** command has not been processed and at least one email recipient using the email event and inventory notification facility has not been set up to receive inventory email notifications. This command also fails if the email infrastructure has not been set up.

An invocation example

In the following example, you send an inventory email notification to all email recipients who are enabled to receive them:

```
sendinventoryemail
```

The resulting output:

No feedback

setemail (Discontinued)

Attention: The **setemail** command is discontinued. E-mail notification can be configured using the following commands: **mkemailserver**, **chemailserver**, **rmemailserver**, **chemail**, and **lsemailserver**.

startemail

Use the **startemail** command to activate the email and inventory notification function. There are no parameters for this command.

Syntax

▶▶ startemail —————▶▶

Parameters

There are no parameters for this command.

Description

This command enables the email event notification service. No emails are sent to users until the **startemail** command has been run and at least one user has been defined to the system.

Note: This command fails if the **chemail** command has not been used to provide adequate configuration details. The following **chemail** parameters must be specified:

- **reply**
- **contact**
- **primary**
- **location**

An invocation example to start the email error notification service

```
startemail
```

The resulting output:

No feedback

stopemail

Use the **stopemail** command to stop the email and inventory notification function. There are no parameters for this command.

Syntax

►► — stopemail ————— ◀◀

Parameters

There are no parameters for this command.

Description

This command stops the email error notification function. No emails are sent to users until the **startemail** command is reissued.

An invocation example to stop the email and inventory notification function

```
stopemail
```

The resulting output:

No feedback

testemail

Use the **testemail** command to send an email notification to one user of the email notification function or to all users of the email notification function to ensure that the function is operating correctly.

Syntax

►► — testemail ————— ◀◀

user_name
user_id
-all

Parameters

user_id | *user_name*

(Required if you do not specify **-all**) Specifies the user ID or user name of the email recipient that you want to send a test email 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 *user_name* or *user_id*) Sends a test email to all email users configured to receive notification of events of any notification type. No attempt is made to send the test email to a user who does not have any notification setting set to *on*.

Description

This command sends test emails to the specified users. The email recipient expects to receive the test email within a specified service time. If the email is not received within the expected time period, the recipient must contact the administrator to ensure that the email settings for the user are correct. If there is still a problem, check your product support information.

The email recipient uses the test email to check that the Simple Mail Transfer Protocol (SMTP) name, the IP address, the SMTP port, and the user address are valid.

An invocation example that sends a test email to user ID manager2008

```
testemail manager2008
```

The resulting output:

No feedback

Chapter 12. Enclosure commands

Storwize V7000, Flex System V7000 Storage Node, Storwize V3500, and Storwize V3700 only: Enclosure commands capture information that can assist you with managing enclosures.

addcontrolenclosure

Use the **addcontrolenclosure** command to add control enclosures to the clustered system.

Syntax

```
➤➤➤ addcontrolenclosure — — -iogrp — io_grp_id_or_name — — -sernum — enclosure_serial_number ➤➤➤
```

➤

```
└─ -site ──┬─ site_name ──┬─
```

 └─ *site_id* ──┬─

 └─

Parameters

-iogrp *io_grp_id_or_name*

The I/O group in which you want to put the control enclosure.

-sernum *enclosure_serial_number*

The serial number of the control enclosure you want to add.

-site *site_name* | *site_id*

(Optional) Specifies the numeric site ID or site name of the new control enclosure. (The site ID can only be set to a value of 1 or 2.)

Description

Use this command to add a control enclosure to the system.

Note: Transparent cloud tiering can be enabled on a system if every node on the system supports it. If a system supports transparent cloud tiering, you cannot add nodes that do not support it to the system.

An invocation example

```
addcontrolenclosure -iogrp 0 -sernum 2361443
```

The resulting output:

```
Enclosure containing Node, id [x], successfully added
```

An invocation example

```
addcontrolenclosure -iogrp 1 -sernum 1234567 -site site2
```


The resulting output:

```
Enclosure, Node id [2], successfully added
```

chenclosure

Use the **chenclosure** command to modify enclosure properties.

Syntax



Parameters

Note: Optional parameters are mutually exclusive. Exactly one of the optional parameters must be set.

-identify *yes* | *no*

(Optional) Causes the identify LED start or stop flashing.

-managed *yes* | *no*

(Optional) Changes the enclosure to 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

(Required) Specifies the enclosure you want to modify.

Description

Use this command to modify enclosure properties.

An invocation example to change the enclosure ID from 7 to 4

```
chenclosure -id 4 7
```

The resulting output:

No feedback

An invocation example to change enclosure 1 to unmanaged

```
chenclosure -managed no 1
```

The resulting output:

No feedback

An invocation example to make the identify LED on enclosure 1 stop flashing

```
chenclosure -identify no 1
```


The resulting output:

No feedback

chenclosurecanister

Use the **chenclosurecanister** command to modify the properties of an enclosure canister.

Syntax



Note: Exactly one of the optional parameters must be set.

Parameters

```
-excludesasport yes | no
```

You can use the **-force** flag if there are dependent volumes.

-force

Important: Using the **-force** parameter might result in a loss of access to your data. Use it only under the direction of your product support group or representative.

(Optional) Resets the enclosure on the canister.

```
-identify yes | no
```

-port 1 | 2 | 3 | 4

Ports 3 and 4 are for Storwize V5000 only.

Specifies the canister to which you want to apply the change.

Specifies the enclosure for which the canister is a member.

This command enables you to modify the properties of an enclosure canister. You must also designate a port (using the `-port` parameter) when this parameter is used.

An invocation example to exclude SAS port 1 on canister 2 of enclosure 1

```
chenclosurecanister -excludesasport yes -port 1 -canister 2 1
```

The resulting output:

No feedback

An invocation example to make the fault LED flash on canister 1 of enclosure 3

```
chenclosurecanister -identify yes -canister 1 3
```

The resulting output:

No feedback

chenclosedisplaypanel

Use the **chenclosedisplaypanel** command to modify the properties for an enclosure display panel.

Syntax

```
➤ chenclosedisplaypanel — [ -clearswap ] [ -displaypanel displaypanel_id ] —  
➤ enclosure_id —
```

Parameters

-clearswap

(Optional) Specifies that the enclosure display panel swap bit be cleared.

-displaypanel *sem_id*

(Optional) Specifies the display panel ID to change. The value must be a number.

enclosure_id

(Required) Specifies the enclosure ID for the enclosure that contains the display panel. The value must be a number 1 - 99.

Description

This command modifies the properties for an enclosure display panel.

An invocation example

```
chenclosedisplaypanel -clearswap -displaypanel 1 3
```

The resulting output:

No feedback

chenclosurepsu

Use the **chenclosurepsu** command to modify the properties of an enclosure power-supply unit (PSU).

Syntax

```
➤ chenclosurepsu — [ -psu psu_id ] — enclosure_id —
```

Parameters

-psu *psu_id*

Identifies the PSU in the enclosure that the command will apply to.

enclosure_id

Identifies the enclosure that the slot is a member of.

Description

This command enables you to modify the properties of an enclosure PSU.

An invocation example

```
chenclosurepsu -psu 2 2
```

The resulting output:

There is no output if the command is successful.

chenclosuresem

Use the **chenclosuresem** command to modify the properties of an enclosure SEM.

Syntax

```
➤ chenclosuresem — [ -clearswap ] [ -enclosuresemid — sem_id ] — enclosure_id ➤
```

Parameters

-clearswap

(Optional) Specifies that the enclosure SEM swap bit be cleared.

-enclosuresemid *sem_id*

(Optional) Specifies the enclosure SEM ID. The value must be a number 1 - 2.

enclosure_id

(Required) Specifies the enclosure ID for the enclosure that contains the SEM. The value must be a number 1 - 99.

Description

This command modifies the properties of an enclosure SEM.

An invocation example

```
chenclosuresem -clearswap -enclosuresemid 1 8
```


The resulting output:

No feedback

chenclosureslot

Use the **chenclosureslot** command to modify the properties of an enclosure slot.

Syntax



Note:

1. Optional parameters are mutually exclusive.
2. You can only specify **-port** or **-force** when you also specify **-exclude**.
3. Exactly one of the optional parameters must be set.
4. Using **-force** has an effect on the operation of **-exclude yes**.

Parameters

-identify *yes* | *no*

(Optional) Change the state of fault light-emitting diode (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 enclosureid**: The port you specify with *port_id* is 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 checks for dependent volumes. If issuing this command would result in losing access to data, then the command fails and an error message displays. 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 enclosureid**: The port is 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 has 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 enclosureid**: If you issue this command without defining a port, the command simultaneously acts on both ports.

-port *1* | *2*

(Optional) Specifies the port on the canister to be excluded. If it is not specified, **-exclude** acts on both ports.

-force

(Optional) Forces the port on the canister to be excluded.

Important: Using the **-force** parameter might result in a loss of access. Use it only under the direction of your product support information.

-slot *slot_id*

(Required) Specifies the slot ID. The value must be a number from 1 - 92.

The slots are numbered 1 (leftmost) to 24 (rightmost) when viewed from in front of the enclosure, in 24-slot enclosures. In 12-slot enclosures, the slots are arranged in numerical order in three rows with four slots. For example, the:

- First row contains slots 1, 2, 3, and 4 (in that order)

- Second row contains slots 5, 6, 7, and 8 (in that order)
- Third row contains slots 9, 10, 11, and 12 (in that order)

enclosure_id

(Required) The enclosure that the slot is a member of.

Description

These commands enable you to modify the properties of an enclosure slot.

An invocation example to turn on the identify LED on slot 7 in enclosure 1

```
chenclosureslot -identify yes -slot 7 1
```

The resulting output:

No feedback

An invocation example to force the exclusion of port 1 of slot 7 in enclosure 3

```
chenclosureslot -exclude yes -port 1 -force -slot 7 3
```

The resulting output:

No feedback

(satask) chenclosurevpd (Deprecated)

The **chenclosurevpd** command is deprecated. Use the **chvpd** command instead.

lsenclosure

Use the **lsenclosure** command to view a summary of the enclosures.

Syntax

```

▶▶ lsenclosure — [ -nohdr ] [ -filtervalue — attribute_value ] [ -filtervalue? ]
▶ [ -delim — delimiter ] [ enclosure_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.

-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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.

- When you use a wildcard, enclose the filter entry within double quotation marks (""): `lsenclosure -filtervalue id="1*"`

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- drive_slots
- id
- IO_group_id
- IO_group_name
- managed
- online_canisters
- online_PSUs
- product_MTM
- serial_number
- status
- total_canisters
- total_PSUs
- type

-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.

enclosure_id

Detailed information for the enclosure that you specify.

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 63 shows the possible outputs:

Table 63. lsenclosure output

Attribute	Description
id	Indicates the ID of the enclosure.
status	Indicates if an enclosure is visible to the SAS network: <ul style="list-style-type: none"> • online if a managed or unmanaged enclosure is visible • offline if 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	Indicates the type of enclosure: <ul style="list-style-type: none"> • control • expansion
managed	Indicates if the enclosure is managed: <ul style="list-style-type: none"> • yes • no

Table 63. *lsenclosure* output (continued)

Attribute	Description
IO_group_id	Indicates the I/O group the enclosure belongs to; blank if canisters are connected to two different I/O groups.
IO_group_name	Indicates the I/O group the enclosure belongs to; blank if canisters are connected to two different I/O groups.
fault_LED	Indicates the status of the fault light-emitting diode (LED) on the enclosure: <ul style="list-style-type: none"> • on if 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 if there is insufficient battery power to run I/O • off if there are faults on the enclosure or its components
identify_LED	Indicates the state of the identify LED: <ul style="list-style-type: none"> • off if the enclosure is not identified • slow_flashing if the enclosure is being identified
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	Indicates the product machine type and model.
serial_number	Indicates 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	Indicates the FRU part number of the enclosure.
FRU_identity	Indicates the 11S serial number that combines the manufacturing part number and the serial number.
total_canisters	Indicates the maximum number of canisters for this enclosure type.
online_canisters	Indicates the number of canisters contained in this enclosure that are online.
total_PSUs	Indicates the number of power and cooling units in this enclosure.
online_PSUs	Indicates the number of power-supply units (PSUs) contained in this enclosure that are online.
drive_slots	Indicates the number of drive slots in the enclosure.
firmware_level_1	Indicates the version of the microcode image (midplane firmware version) installed on the midplane.
firmware_level_2	Indicates the version of the midplane metadata (midplane vital product data, or VPD, version) installed on the midplane.
machine_part_number	Blank.
machine_signature	Indicates a machine signature unique to the control enclosure and representing the serial number and machine part number. The format is a hyphenated string of nineteen hexadecimal characters. Remember: Expansion enclosures do not have a machine signature.
interface_speed	Indicates the SAS interface speed of the enclosure (in gigabits per second, or Gbps). The values are: <ul style="list-style-type: none"> • 6 Gbps • 12 Gbps • Blank for an unknown or unsupported enclosure
total_sems	Indicates the total number of secondary expander modules (SEMs) that are in the system. The value must be a number 0 - 2.
online_sems	Indicates the total number of SEMs in the system that are online. The value must be a number 0 - 2.

A detailed invocation example

`lsenclosure 1`

The detailed resulting output:

```
id 1
status online
type control
managed no
IO_group_id 0
IO_group_name io_grp0
fault_LED off
identify_LED off
error_sequence_number
product_MTM 2072-02A
serial_number 64G005S
FRU_part_number 85Y5896
FRU_identity 11S85Y5962YHU9994G005S
total_canisters 2
online_canisters 2
total_PSUs 2
online_PSUs 2
drive_slots 12
firmware_level_1 10
firmware_level_2 F6C07926
machine_part_number 2072L2C
machine_signature 0123-4567-89AB-CDEF
ambient_temperature 30
total_fan_modules:2
online_fan_modules:2
interface_speed:6Gb
total_sems 2
online_sems 1
```

lsenclosurebattery

Use the **lsenclosurebattery** command to display information about the batteries. The batteries are located in the node canisters.

Syntax

```
lsenclosurebattery [-nohdr] [-filtervalue attribute_value]
                  [-filtervalue?] [-delim delimiter]
                  [-battery battery_id] enclosure_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.

-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 when using the 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:

```
lsenclosurebattery -filtervalue "battery_id=1"
```

-filtervalue?

(Optional) Displays a list of valid filter attributes for the **-filtervalue** *attribute=value* parameter:

- battery_id
- charging_status
- enclosure_id
- end_of_life_warning
- percent_charged
- recondition_needed
- status

-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 width of each item of data. A detailed view provides each item of data in 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.

-battery *battery_id*

(Optional) Provides a detailed view of the specified battery. Valid only when an enclosure is specified.

enclosure_id

(Optional) Lists the batteries for the specified enclosure.

Description

This command displays information about the batteries, which are located in the node canisters. The concise view displays a line for each battery slot in every enclosure whether a battery exists for that slot or not. Batteries are not shown for expansion enclosures. Table 64 shows possible outputs.

Table 64. *lsenclosurebattery* outputs

Attribute	Description
enclosure_id	Identifies the identity of the enclosure that contains the battery.
battery_id	Identifies the battery in the enclosure.
status	Identifies the status of the battery: <ul style="list-style-type: none">• online indicates the battery is present and working as usual.• degraded indicates the battery is present but not working as usual.• offline indicates the battery cannot be detected.

Table 64. *lsenclosurebattery* outputs (continued)

Attribute	Description
charging_status	Identifies the charging state of the battery: <ul style="list-style-type: none"> charged indicates the battery is fully charged. charging indicates the battery is charging. discharging indicates the battery is losing voltage (life). idle indicates the battery is neither charging nor discharging but it is not fully charged.
percent_charged	Indicates the charge of battery (in a percentage).
end_of_life_warning	Identifies the battery's end of life (with a warning noise). The values are yes and no. Important: Replace the battery.
FRU_part_number	Identifies the FRU part number of the battery.
FRU_identity	Identifies the 11S number, combining the manufacturing part number and the serial number.
firmware_level	The version of the microcode image that is installed on the battery.
error_sequence_number	Indicates the error log (or event log) number of the highest priority error for this object. This output field 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.
remaining_charge_capacity_mAh	Identifies the battery's remaining charge capacity in milliampere-hour (mAh).
full_charge_capacity_mAh	Identifies the fully-charged capacity of the battery in mAh (this diminishes as the battery ages).
compatibility_level	Identifies the battery driver software must support this level to operate with this battery - this comes from the battery vital product data (VPD).
last_recondition_timestamp	Identifies a system timestamp, in the format YYMMDDHHMMSS, for when the last successful calibration of the gas gauge occurred, where: <ul style="list-style-type: none"> Y is year (The first) M is month D is day H is hour (The second) M is minute S is second
powered_on_hours	Identifies the number of hours the battery has been in a powered node (not necessarily the same node).
cycle_count	Identifies the number of charge or discharge cycles performed on the battery.

A concise invocation example

```
lsenclosurebattery 1
```

The resulting output:

```
enclosure_id battery_id status charging_status recondition_needed percent_charged end_of_life_warning
1             1         online idle                no                100              no
1             2         online idle                no                100              no
```

A detailed invocation example

```
lsenclosurebattery -battery 1 1
```

The resulting output:

```

enclosure_id 1
battery_id 1
status online
charging_status idle
recondition_needed no
percent_charged 100
end_of_life_warning no
FRU_part_number 31P1807
FRU_identity 11S00AR085YM30BG42R04P
firmware_level 105:1
error_sequence_number
remaining_charge_capacity_mAh 3477
full_charge_capacity_mAh 3795
compatibility_level 1
last_recondition_timestamp 140528045617
powered_on_hours 1162
cycle_count 10

```

Iscontrolenclosurecandidate (Storwize family products only)

Use the **Iscontrolenclosurecandidate** command to display a list of all control enclosures you can add to the current system.

Syntax

```

>> Iscontrolenclosurecandidate — [ -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

Table 65 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 65. Iscontrolenclosurecandidate attribute values

Attribute	Value
serial_number	Indicates the serial number for the enclosure.
product_MTM	Indicates the MTM for the enclosure.
machine_signature	Indicates the machine signature for the enclosure.

A concise invocation example

lscontrolenclosurecandidate

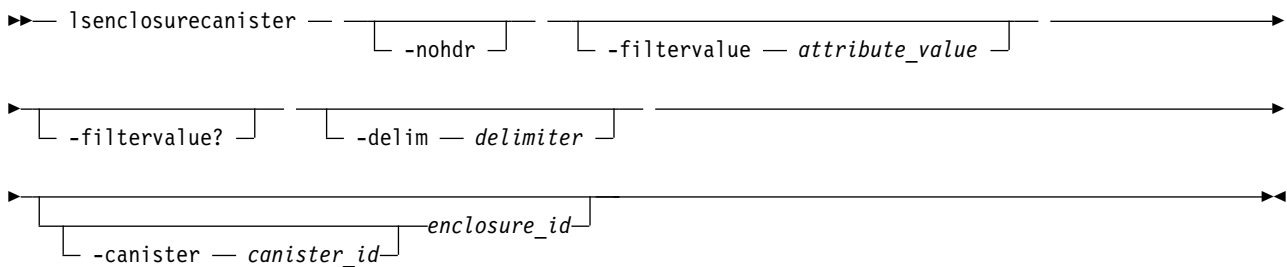
The concise resulting output:

```
serial_number product_MTM machine_signature
G00F7GY       2076-124 5746-9812-B5CF-FEF9
```

lsenclosurecanister

Use the **lsenclosurecanister** command to view a detailed status for each canister in an enclosure.

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.

-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 when using the 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:

```
lsenclosurecanister -filtervalue "node_name=node*"
```

-filtervalue?

(Optional) Displays a list of valid filter attributes for the **-filtervalue attribute=value** parameter:

- enclosure_id
- canister_id
- node_id
- node_name
- status
- type

-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.

-canister *canister_id*

Valid only when the *enclosure_id* is specified. Provides a detailed view of the canister for the specified enclosure.

enclosure_id

Lists the canisters for the specified enclosure.

Description

This command enables you to view a detailed status for each canister in an enclosure. Table 66 shows the possible outputs:

Table 66. Isenclosurecanister 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 offline or not part of the clustered system.
<i>node_name</i>	The node that corresponds to this canister; blank if the canister is not a node, or if the node is offline or not part of the clustered system.
<i>temperature</i>	The temperature of the canister in degrees centigrade.
<i>identify_LED</i>	The status of the <i>identify_LED</i> . The values are on, off, or slow-flashing.
<i>fault_LED</i>	The status of the <i>fault_LED</i> . The values are on, off, or slow-flashing.
<i>SES_status</i>	The SCSI status of the connection between the device and the canister. The values are online and offline.
<i>FRU_part_number</i>	The field-replaceable unit (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 worldwide node name (WWNN) of the canister (node canisters only).
<i>temperature</i>	(0 to 245) The temperature of the canister (in degrees Celsius). If the temperature goes below 0, s0 is displayed. Remember: The temperature value is not an ambient temperature value. It is an internal temperature sensor value.

Table 66. *lsenclosurecanister output (continued)*

Attribute	Description
fault_LED	The state of the combined fault and identify light-emitting diodes (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
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 (meaning logged in but cannot communicate with the canister) • degraded (meaning the SAS cable is not fully functional) • Blank (which can appear on control canisters; see lsportsas)
SAS_port_2_status	Indicates if there is damage to the cable between SAS ports: <ul style="list-style-type: none"> • online • offline • excluded (meaning logged in but cannot communicate with the canister) • degraded (meaning the SAS cable is not fully functional) • Blank (which can appear on control canisters; see lsportsas)
firmware_level	The firmware level of the Small Computer System Interface (SCSI) Enclosure Services (SES) code, or canister firmware version, running on the canister.
firmware_level_2	The version of the first other microcode image (canister bootloader version) installed on the canister.
firmware_level_3	The version of the second other microcode (canister complex programmable logic device, or CPLD, version) image installed on the canister.
firmware_level_4	The version of the third other microcode image (canister flash configuration version) installed on the canister.
firmware_level_5	The version of the canister metadata (canister VPD version) installed on the canister.

A concise invocation example

```
lsenclosurecanister -delim :
```

The resulting output:

```
enclosure_id:canister_id:status:type:node_id:node_name
1:1:degraded:expansion:1:node1
```

A detailed invocation example

```
lsenclosurecanister -canister 1 1
```

The detailed 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

firmware_level_2 0501
firmware_level_3 14
firmware_level_4 B69F66FF
firmware_level_5 5C2A6A44

```

lsenclosurechassis (Flex V7000 Storage Node)

Use the **lsenclosurechassis** command to provide a description of the chassis-specific enclosure properties, including its location within the chassis.

Syntax

```

➤ — lsenclosurechassis — [ -nohdr ] [ -delim delimiter ] [ enclosure_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.

enclosure_id

(Optional) Indicates the unique enclosure identifier (a number between 1 and 99).

Description

This command has both a detailed and concise view. The *enclosure_id* keyword is required for the detailed view.

The following table displays information about chassis-specific enclosure properties and shows possible outputs.

Table 67. *lsenclosurechassis* outputs

Attribute	Description
enclosure_id	Specifies the enclosure identifier. It is a numeric character between the numbers 1 and 99.
chassis_name	Specifies the chassis name. It can be set from the CMM, and is blank or an alphanumeric string containing up to 128 characters.
canister_1_bay	Specifies the first canister bay's enclosure position within the chassis. It is a numeric character between the numbers 0 and 254.
canister_2_bay	Specifies the second canister bay's enclosure position within the chassis. It is a numeric character between the numbers 0 and 254.
numbering_scheme	Specifies the chassis numbering scheme set from the CMM. It can be a numeric character between the numbers 0 and 255.
pos_in_rack	Specifies the chassis position within the rack set from the CMM. It must be an alphanumeric 2-character string.
rack_location	Specifies the location of the rack containing the chassis set from the CMM. It can be blank or an alphanumeric string containing up to 128 characters.
rack_room	Specifies the room that contains the rack set from the CMM. It can be blank or an alphanumeric string containing up to 128 characters.
chassis_mtm	Specifies the chassis machine type or model. The type or model is an alphanumeric string containing up to 22 characters.
chassis_sn	Specifies the chassis serial number. The serial number is an alphanumeric string containing up to 22 characters.
chassis_uuid	Specifies the chassis unique user identifier. The identifier is an alphanumeric string containing up to 128 characters.
chassis_rack	Specifies the identifier for the rack that contains the chassis. The identifier is blank or an alphanumeric string containing up to 128 characters.

An invocation example

```
lsenclosurechassis 1
```

The resulting output:


```
enclosure_id 1
chassis_name 25631
  canister_1_bay 0
  canister_2_bay 0
numbering_scheme 0
pos_in_rack 1
rack_location In the corner
rack_room D-East
chassis_mtm 2078-219
chassis_sn 64H123R
chassis_uuid 987654321
chassis_rack Rack47
```

lsenclosedisplaypanel

Use the **lsenclosedisplaypanel** command to display information about the display panel in an enclosure.

Syntax

```
➤— lsenclosedisplaypanel — -displaypanel — displaypanel_id — ————→
                                     └─ -nohdr ─┘
```



Parameters

-displaypanel *displaypanel_id*

(Required) Specifies the display panel ID for the display panel that is being displayed. The value must be a number.

-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 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 1-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.

enclosure_id

(Required) Specifies the enclosure ID for the enclosure data that is being displayed. The value must be a number 1 - 99.

Description

This command displays information about the display panel in an enclosure.

This table provides the attribute values that can be displayed as output view data.

Table 68. ? output

Attribute	Description
enclosure_id	Indicates the enclosure ID for the enclosure that contains the display panel. The value must be a number 1 - 99.
displaypanel_id	Indicates the ID of the display panel that is in the enclosure. The value must be a number.
status	Indicates the display panel status for the display panel that is in the enclosure. The values are: <ul style="list-style-type: none"> enum online degraded offline
error_sequence_number	Indicates the event log sequence number for the current event that is logged against the secondary expander module (SEM). The value is blank if there is no event to log.
FRU_part_number	Indicates the FRU part number of the display panel. The value must be a 7-character numeric string.
FRU_identity	Indicates the FRU identity of the display panel. The value must be a 22-character alphanumeric string.

A concise invocation example

```
lsenclosuredisplaypanel
```

The resulting output:

enclosure_id	display_panel_id	status
1	1	online
2	1	online
3	1	online

A detailed invocation example

```
lsenclosuredisplaypanel -displaypanel 1 3
```


The resulting output:

```
enclosure_id 3
displaypanel_id 1
status online
error_sequence_number
FRU_Part_Number *****
FRU_Identity *****
```

lsenclosurefanmodule

Use the **lsenclosurefanmodule** command to report the status of each fan module and the contained fans in an enclosure.

Syntax



Parameters

-fanmodule fanmodule_id

(Optional) Specifies the ID of the fan module for which data is displayed. The possible values are 1 or 2, and any other value returns no output.

-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.

enclosure_id

(Optional) Specifies the ID of the enclosure for which data is displayed.

Description

The command reports status for fan modules and the contained fans in an enclosure.

Table 69 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 69. Isenclosurefanmodule attribute values

Attribute	Value
enclosure_id	Indicates the enclosure ID of the enclosure that contains the fan module.
fan_module_id	Indicates the fan module ID of the fan module that is in the enclosure. The possible values are 1 or 2.
status	Indicates the combined status of the fan module and any contained fans. The values are: <ul style="list-style-type: none">• online• offline• degraded
error_sequence_number	Indicates the event log sequence number of the current event that is logged against the fan module. It is blank if there is no current event.
FRU_part_number	Indicates the part number of the fan module.
FRU_identity	Indicates the FRU identity of the fan module.
fault_LED	Indicates the status of the fault light-emitting diode (LED) on the fan module: <ul style="list-style-type: none">• on, which indicates that the LED is on• off, which indicates that the LED is off• unknown, which indicates that the LED status is unknown

An invocation example

```
lsenclosurefanmodule
```

The resulting output:

```
enclosure_id fan_module_id status
1             1             online
1             2             online
2             1             online
2             2             online
```

An invocation example

```
lsenclosurefanmodule 2
```

The resulting output:

```
enclosure_id fan_module_id status
2             1             online
2             2             online
```

An invocation example

```
lsenclosurefanmodule -fanmodule 1 1
```

The resulting output:

```
enclosure_id 1
fan_module_id 1
status online
```

```
error_sequence_number
FRU_part_number 31P1847
FRU_identity 11S31P1846YM10BG3B101N
fault_LED off
```

lsenclosurepsu

Use the **lsenclosurepsu** command to view information about each power-supply unit (PSU) in the enclosure.

Syntax

```
lsenclosurepsu [-nohdr] [-filtervalue attribute_value]
                [-filtervalue?] [-delim delimiter] [-psu psu_id] enclosure_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.

-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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):

```
lsenclosurepsu -filtervalue "psu_id=1"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- enclosure_id
- psu_id
- status

-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.

-psu *psu_id*

(Optional) Valid only when the **enclosure_id** is specified. Provides a detailed view of the PSU for the specified enclosure.

enclosure_id

(Optional) Lists the PSUs for the specified enclosure.

Description

This command enables you to view information about each power-supply unit (PSU) in the enclosure. Table 70 shows the possible outputs:

Table 70. Isenclosurepsu output

Attribute	Description
enclosure_id	Indicates the ID of the enclosure containing the PSU.
psu_id	Indicates the ID of the PSU in the enclosure.
status	Indicates the status of the power and cooling unit in the enclosure: <ul style="list-style-type: none">• online indicates a PSU is present and working normally• offline indicates a PSU cannot be detected• degraded indicates a PSU is present but not working normally
input_failed	<ul style="list-style-type: none">• on indicates no usable input power is detected from the power distribution unit.• off indicates the input power is OK.
output_failed	<ul style="list-style-type: none">• on indicates no usable output power is detected from the distribution unit.• off indicates the output power is OK.
input_power	<ul style="list-style-type: none">• ac indicates the power supply requires AC power input.• dc indicates the power supply requires DC power input.• unknown indicates the power supply is not known or cannot be determined.
fan_failed	<ul style="list-style-type: none">• on indicates that if the AC, DC, and fan LEDs are all on, there is a PSU fault. If only the fan LED is on, then there is a fan failure.• off indicates the fans in this PSU are OK.
redundant	Indicates (yes or no) whether you can remove the power supply: <ul style="list-style-type: none">• If the PSU is on an expansion enclosure, the other PSU must be online.• If the PSU is on a control enclosure, 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	Indicates the FRU part number of the PSU.
FRU_identity	Indicates the 11S number, combining the manufacturing part number and the serial number.
firmware_level_1	Indicates the version of the microcode image (power supply firmware version) installed on the power supply.
firmware_level_2	Indicates the version of the power supply metadata (power supply vital product data, or VPD, version) installed on the power supply. Note: This field might not be applicable for some systems and is blank for all PSU types.
firmware_level_3	Indicates the version of the secondary microcode image installed on the enclosure's High Efficiency (HE) power supply unit (PSU). Note: This field might not be applicable for some systems and is blank for all PSU types.

An invocation example

```
lsenclosurepsu -delim :
```

The resulting output:

```
enclosure_id:PSU_id:status:input_power
1:1:online:ac
1:2:online:ac
```

A detailed invocation example

```
lsenclosurepsu -psu 1 1
```

The detailed resulting output:

```
enclosure_id 1
PSU_id 1
status online
input_failed off
output_failed on
fan_failed off
redundant yes
error_sequence_number
FRU_part_number 85Y5847
FRU_identity 11S85Y5847YG50CG07W0LJ
firmware_level_1 0314
firmware_level_2 AF9293E5
firmware_level_3
input_power ac
```

lsenclosuresem

Use the **lsenclosuresem** command to display the status (or any pertinent data) about secondary expander modules (SEM) in a 5U92 system.

Syntax

```
lsenclosuresem — [ -sem sem_id ] [ -nohdr ] [ -delim delimiter ]
enclosure_id
```

Parameters

-sem *sem_id*

(Optional) Specifies the SEM ID for the SEM data that is being 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 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 1-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.

enclosure_id

(Required) Specifies the enclosure ID for the enclosure data that is being displayed. The value must be an integer 1 - 99.

Description

This command displays the status (or any pertinent data) about SEM in a 5U92 system.

A 5U92 system is a 5U enclosure that can contain up to 92 3.5 inch drives (but can only be used in an expansion enclosure).

This table provides the attribute values that can be displayed as output view data.

Table 71. lsenclosuresem output

Attribute	Description
id	Indicates the enclosure ID of the enclosure that contains the disk drawer. The value must be a number 1 - 99.
sem_id	Indicates the SEM ID for the SEM that is in the enclosure. The value must be a number, 1 or 2.
status	Indicates the SEM status for the SEM that is in the enclosure. The values are: <ul style="list-style-type: none"> • online, which indicates that the SEM is online • degraded, which indicates that the SEM is degraded • offline, which indicates that the SEM is offline
expander_1_status	Indicates the status of the first or lowest order expander index. The values are: <ul style="list-style-type: none"> • online, which indicates that the SEM is online • degraded, which indicates that the SEM is degraded • offline, which indicates that the SEM is offline
expander_2_status	Indicates the status of the second expander index (or $e+1$, where e is the lowest-order or first expander index). The values are: <ul style="list-style-type: none"> • online, which indicates that the SEM is online • degraded, which indicates that the SEM is degraded • offline, which indicates that the SEM is offline
error_sequence_number	Indicates the event log sequence number of the current event that is logged against the SEM. The value is blank if there is no error to log.
FRU_part_number	Indicates the field-replaceable unit (FRU) part number for the SEM. The value must be a 7-character numeric string.
FRU_identity	Indicates the FRU ID for the SEM. The value must be a 22-character alphanumeric string.
firmware_level_1	Indicates the SCSI Enclosure Services (SES) firmware level of the lowest order expander index. The value must be a 22-character alphanumeric string.
firmware_level_2	Indicates the bootloader firmware level of the lowest order expander index. The value must be a 22-character alphanumeric string.
firmware_level_3	Indicates the SES firmware level of the second lowest order expander index. (or $s+1$, where s is the lowest-order or first expander index). The value must be a 22-character alphanumeric string.
firmware_level_4	Indicates the bootloader firmware level of the second lowest order expander index. (or $b+1$, where b is the lowest-order or first expander index). The value must be a 22-character alphanumeric string.

Table 71. *lsenclosuresem* output (continued)

Attribute	Description
firmware_level_5	Indicates the complex programmable logic device (CPLD) firmware level of the lowest order expander index. The value must be a 22-character alphanumeric string.
firmware_level_6	Indicates the CPLD firmware level of the second lowest order expander index (or <i>c+1</i> , where <i>c</i> is the lowest-order or first expander index). The value must be a 22-character alphanumeric string.

A concise invocation example

```
lsenclosuresem 1
```

The resulting output:

```
enclosure_id sem_id status expander1_status expander2_status
1           1      online online           online
1           2      online online           online
```

A detailed invocation example

```
lsenclosuresem -sem 1 1
```

The resulting output:

```
enclosure_id 1
sem_id 1
status online
expander1_status online
expander2_status online
error_sequence_number
FRU_Part_Number *****
FRU_Identity *****
firmware_level_1 0802.official
firmware_level_2 000E
firmware_level_3 0802.official
firmware_level_4 000E
firmware_level_5 1A.04.E3
firmware_level_6 1A.04.E5
```

lsenclosureslot

Use the **lsenclosureslot** command to view information about each drive slot in the enclosure.

Syntax

```

>>> lsenclosureslot — [ -filtervalue — attribute_value ] [ -filtervalue? ]
> [ -delim — delimiter ] [ -nohdr ] [ -slot — slot_id ] enclosure_id

```

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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):


```
lsenclosureslot -filtervalue "enclosure_id>2"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- drive_id
- drive_present
- enclosure_id
- port_1_status
- port_2_status
- slot_id

-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.

-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. This parameter suppresses the display of these headings.

-slot slot_id

(Optional) Specifies the slot to display information for (it provides a detailed view for that enclosure slot). This parameter is valid only when an enclosure is specified. The value must be a number from 1 - 92.

Note:

If slot information is requested for a slot that does not exist on the specified enclosure, value displayed is blank.

enclosure_id

(Optional) Lists slots for that enclosure. Must be specified if **-slot** is used.

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 72 shows the possible outputs:

Table 72. lsenclosureslot output

Attribute	Description
enclosure_id	The identity of the enclosure which contains the drive slot.
slot_id	Identifies which of the drive slots in the enclosure this is.

Table 72. *lsenclosureslot* output (continued)

Attribute	Description
port_1_status	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"> • online: enclosure slot port 1 is online • excluded_by_drive: the drive excluded the port • excluded_by_enclosure: the enclosure excluded the port • excluded_by_system: the clustered system (system) has excluded the port
port_2_status	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"> • online: enclosure slot port 2 is online • excluded_by_drive: the drive excluded the port • excluded_by_enclosure: the enclosure excluded the port • excluded_by_system: the clustered system (system) has excluded the port
fault_LED	The state of the combined fault and identify light-emitting diodes (LEDs): <ul style="list-style-type: none"> • off, which indicates no fault • slow_flashing, which indicates the identify mode <p>Note: When the LED is in identify mode, it conceals whether there is a fault present (it always flashes). When you remove it from identity mode, the LED is turned on or off.</p> <ul style="list-style-type: none"> • on: fault
powered	Indicates whether the slot is powered on. <ul style="list-style-type: none"> • yes • no
drive_present	Indicates if a drive is in the slot. The drive can be working, dead, or powered off. <ul style="list-style-type: none"> • yes (present) • no (empty)
drive_id	Indicates 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.
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.
interface_speed	Indicates the lowest interface speed for the connected drive slot (in gigabits per second, or Gbps). The values are: <ul style="list-style-type: none"> • 1.5 Gbps • 3 Gbps • 6 Gbps • 12 Gbps • Blank if both ports are isolated or the drive is not connected
row	Identifies the row in which the slot appears. The values are A - G.
column	Identifies the column in which the slot appears. The values are 1 - 14.

A concise invocation example

This example displays information about mappings between the 1 and 2 dimensional IDs

```
lsenclosureslot
```

The resulting output:

enclosure_id	slot_id	port_1_status	port_2_status	drive_present	drive_id	row	column
1	1	online	online	no		A	1
1	2	online	online	no		A	2
1	3	online	online	no		A	3
1	4	online	online	no		A	4
1	5	online	online	no		A	5
...							
1	87	online	online	no		G	9
1	88	online	online	no		G	10
1	89	online	online	no		G	11
1	90	online	online	no		G	12
1	91	online	online	no		G	13
1	92	online	online	no		G	14

A detailed invocation example showing slot 2 in enclosure 5

```
lsenclosureslot -delim : -slot 2 5
```

The resulting output:

```
enclosure_id:5
slot_id:2
port_1_status:online
port_2_status:online
fault_LED:off
powered:yes
drive_present:yes
drive_id:105
error_sequence_number:
interface_speed:6Gb
```

A detailed invocation example

```
lsenclosureslot -delim :
```

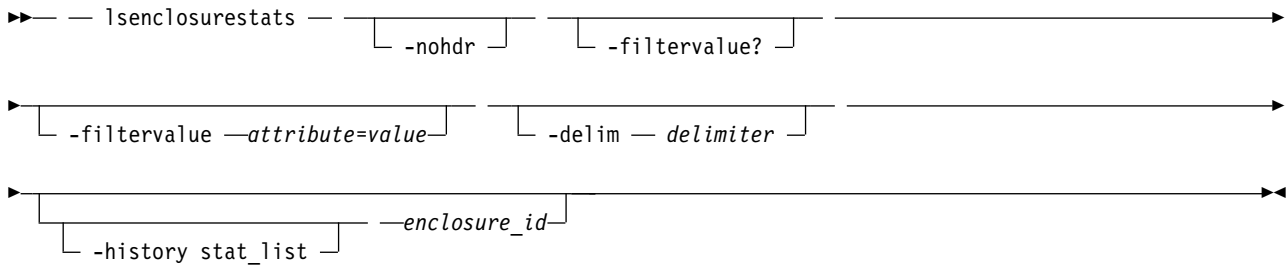
The resulting output:

```
enclosure_id:slot_id:port_1_status:port_2_status:drive_present:drive_id:error_sequence_number
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:
```

lsenclosurestats

Use the **lsenclosurestats** command to display the most recent values (averaged) of all enclosure statistics. It can also display a history of those values for any given subset of the available statistics.

Syntax



Parameters

-history stat_list

(Optional) Produces a history of values for enclosure statistics.

-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.

-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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):

```
lsenclosurestats -filtervalue "enclosure_id>2"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- enclosure_id
- stat_name

-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.

enclosure_id

(Optional) Indicates the unique enclosure identifier (a number between 1 and 99).

Description

Remember: This command cannot be used for products that do not support environmental statistics.

If you specify **-history stat_list** you must also specify *enclosure_id*. Filtering is supported for the concise view but not the detailed view.

Multiple statistical histories can be requested. The limit is the current maximum number of different statistical names published in the concise view. The concise view defines the output order.

For the detailed view, enclosure power is averaged over thirty seconds to provide immediate power.

Note: Averaging applies only to populated samples
Enclosure power is not averaged in the output if **-history** is specified.

This is an invocation example for products that do not support environmental statistics - a message is displayed:

```
lsenclosurestats
```

This is the resulting output:

```
CMMVC6051E An unsupported action was selected.
```

Table 73 displays information about chassis-specific enclosure properties and shows possible outputs for products that support environmental statistics.

Table 73. lsenclosurestats outputs

Attribute	Description
enclosure_id	Indicates the enclosure identifier; it can be a numeric character between 1 and 264.
sample_time	Indicates the time during which the sample occurred.
stat_name	Indicates the name of the statistical field.
stat_current	Indicates the current value of the statistical field.
stat_peak	Indicates the peak value of the statistic field. The last five minutes is used for samples.
stat_peak_time	Indicates the time that the peak occurred.
stat_value	Indicates the value of the statistic.

Remember: Filtering is supported on the *enclosure_id* and *stat_name* fields using the concise view.

An invocation example

```
lsenclosurestats
```

The resulting output:

```
enclosure_id stat_name stat_current stat_peak stat_peak_time
1            power_w  2200          2500    120402103212
1            temp_c   35           36      120402103212
1            temp_f   95           97      120402103212
2            power_w  2300          2600    120402102917
2            temp_c   36           37      120402102917
2            temp_f   97           98      120402102917
4            power_w  2100          2400    120402103202
4            temp_c   33           35      120402103202
4            temp_f   93           95      120402103202
```

An invocation example

```
lsenclosurestats -history power_w 1
```

The resulting output:

enclosure_id	sample_time	stat_name	stat_value
1	120402105137	power_w	2282
1	120402105142	power_w	2290
1	120402105147	power_w	2281
1	120402105152	power_w	2290
1	120402105157	power_w	2281
1	120402105202	power_w	2289
1	120402105207	power_w	2282
1	120402105212	power_w	2289
1	120402105217	power_w	2281
1	120402105222	power_w	2289
1	120402105227	power_w	2281
1	120402105232	power_w	2290
1	120402105237	power_w	2282
1	120402105242	power_w	2289
1	120402105247	power_w	2282
1	120402105252	power_w	2289
1	120402105257	power_w	2282
1	120402105302	power_w	2289
1	120402105307	power_w	2282
1	120402105312	power_w	2289
1	120402105317	power_w	2282
1	120402105322	power_w	2287
1	120402105327	power_w	2281
1	120402105332	power_w	2290
1	120402105337	power_w	2281
1	120402105342	power_w	2289
1	120402105347	power_w	2282
1	120402105352	power_w	2289
1	120402105357	power_w	2281
1	120402105402	power_w	2289
1	120402105407	power_w	2281
1	120402105412	power_w	2289
1	120402105417	power_w	2282
1	120402105422	power_w	2289
1	120402105427	power_w	2282
1	120402105432	power_w	2289
1	120402105437	power_w	2281
1	120402105442	power_w	2290
1	120402105447	power_w	2281
1	120402105452	power_w	2290
1	120402105457	power_w	2282
1	120402105502	power_w	2287
1	120402105507	power_w	2281
1	120402105512	power_w	2290
1	120402105517	power_w	2281
1	120402105522	power_w	2289
1	120402105527	power_w	2282
1	120402105532	power_w	2290
1	120402105537	power_w	2281
1	120402105542	power_w	2290
1	120402105547	power_w	2281
1	120402105552	power_w	2290
1	120402105557	power_w	2281
1	120402105602	power_w	2289
1	120402105607	power_w	2282
1	120402105612	power_w	2289
1	120402105617	power_w	2281
1	120402105622	power_w	2289
1	120402105627	power_w	2281
1	120402105632	power_w	2290

This table provides the possible values that are applicable to the values that are displayed for the **stat_name** attribute.


Table 74. Stat_name field values

Value	Description
power_w	Displays the power consumed in watts.
temp_c	Displays the ambient temperature in Celsius.
temp_f	Displays the ambient temperature in Fahrenheit.

Issasfabric

Use the **Issasfabric** command to see which canisters are visible to a node, and the order of these canisters.

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 displayed.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""): `Issasfabric -filtervalue "status"`

-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?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- enclosure_id
- canister_id
- canister_port_id

- control_enclosure_id
- node_canister_id
- node_canister_port_id
- position
- IO_group_id
- IO_group_name
- node_id
- node_name

Description

Use this command to see which canisters are visible to a node, and the order of these canisters. Table 75 describes possible outputs.

Table 75. 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. If the node does not reside inside a canister or an enclosure this field is blank.
node_canister_id	The identity of the canister the strand comes from. If the node does not reside inside a canister or an enclosure this field is blank.
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.

```
lssasfabric
```

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.

This is the first part of the resulting 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

This is the second part of the resulting output:

node_canister_port_id	position	IO_group_id	IO_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

An invocation example with two enclosures

This examples shows the output when using this command for a pair of expansion enclosures that are wired correctly to a set of nodes.

lssasfabric

The resulting output:

enclosure_id	canister_id	canister_port_id	control_enclosure_id	node_canister_id	node_canister_port_id	position	IO_group_id	IO_group_name	node_id	node_name
1	1	1		1	1	0	0	io_grp0	1	node1
2	1	1		2	1	0	0	io_grp0	1	node1
1	2	1		1	1	0	0	io_grp0	2	node2
2	2	1		2	1	0	0	io_grp0	2	node2

[edit]

resetleds

Use the **resetleds** command to simultaneously switch off all light-emitting diodes (LEDs) in the clustered system (system), including node battery LEDs.

Syntax

▶▶—resetleds—▶▶

Parameters

None.

Description

The **resetleds** command simultaneously switches off all LEDs in the system, including node battery LEDs. This ensures that any identity LED switched on is the only one in the system that is switched on. This command works only on LEDs that are on systems that can communicate, which means they are online or of a supported type. This command fails if an object is offline, or if the enclosure is an unsupported type. This command does not affect LEDs:

- On independently-controlled objects
- On offline objects
- With hardware-only controls

An invocation example

resetleds

The resulting output:

No feedback

triggerenclosuredump

Use the **triggerenclosuredump** command to force the specified enclosure or enclosures to dump data.

Syntax

```
▶▶— triggerenclosuredump — [ -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*

(Optional) 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_id* of either node canister are dumped.

-iogrp *iogrp_id_or_name*

(Optional) The ID or name of the I/O group the control enclosure belongs to.

-enclosure *enclosure_id*

(Optional) The ID of the enclosure you want to dump.

Description

Important: One of the optional parameters must be specified.

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 is one file for each canister successfully dumped and they may be located on different nodes. Dumps are for use by your product support team (or information) if it 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. The system limits the number of enclosure statesaves in the directory to 20 per node.

To trigger enclosure dumps from all enclosures connected to port 1 of the control enclosure in iogrp 2

```
triggerenclosuredump -port 1 -iogrp 2
```

The resulting output:

The data is dumped to the /dumps/enclosure directory if command is successful.

To trigger enclosure dumps from enclosure 5

```
triggerenclosuredump -enclosure 5
```

The resulting output:

The data is dumped to the /dumps/enclosure directory if command is successful.

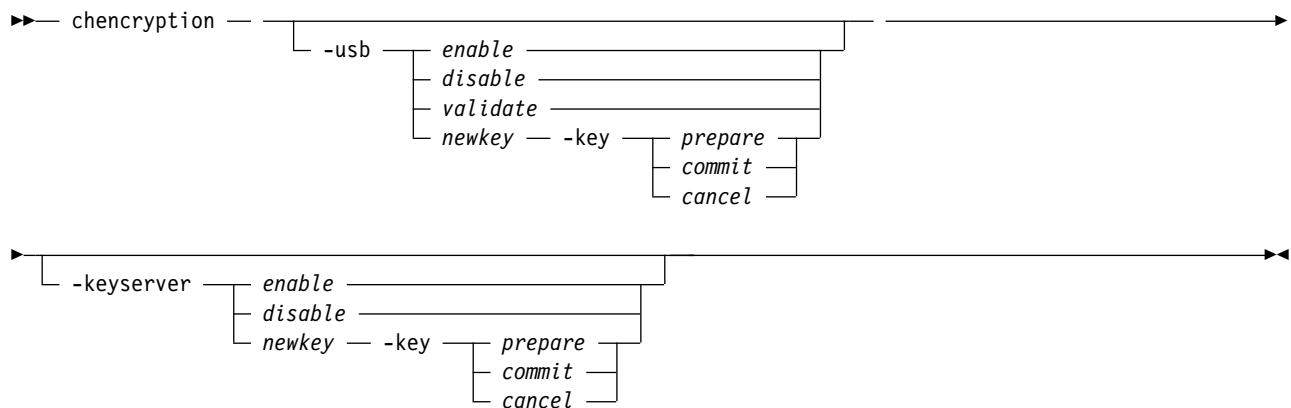
Chapter 13. Encryption commands

Security commands are used to create, change, or list details about clustered system (system) encryption features.

chencryption

Use the **chencryption** command to manage the encryption state of the system.

Syntax



Parameters

This applies to Storwize V7000 and IBM FlashSystem™:

-usb enable | disable | validate | newkey

(Required if you do not specify **-keyserver**) Specifies whether encryption is enabled (or not enabled) or the encryption keys are validated. You can also create new encryption keys which are also stored on universal serial bus (USB) flash drives for use if the system forgets the encryption keys.

-usb enable

Enables encryption capability on the system. Then specify **-usb newkey** to create new keys. Use this command when the system has encryption hardware and encryption licenses (for example, the **lsencryption** value for status is set to licensed).

-usb disable

Disables the encryption capability of the system. If no encryption key is prepared this operation is complete and no further action is needed. Do not use this command if an encryption key is prepared or encrypted arrays already exist.

Remember: This removes all encryption keys (that are not on the USB flash drive) from the system.

-usb validate

Verifies that encryption keys are present on the USB flash drive and makes sure that the keys match the system encryption keys. Use this command when encryption is enabled and encryption keys exist (for example, **lsencryption** value for usb_rekey is set to no).

-usb newkey

Generates a new encryption key on a USB flash drive attached to the system. Use this

command only if the minimum number of USB flash drives that can be used as key material stores are attached to the system (as reported by **lsportusb**). When specifying this parameter, the **-key** option must also be supplied.

-keyserver *enable | disable | newkey*

(Required if you do not specify **-usb**) Specifies the encryption task that involves encryption keys that are managed by key servers.

-keyserver *enable*

Enables encryption capability on the system. Use this command when the system has encryption hardware and encryption licenses (for example, the **lsencryption** value for status is set to **licensed**).

-keyserver *disable*

Disables the encryption capability of the system. If no encryption key is prepared, this operation is complete and no further action is needed. Do not use this command if an encryption key is prepared or encrypted arrays already exist.

-keyserver *newkey*

Generates a new encryption key on the primary key server that is attached the system. You must also specify **-key** when you specify this parameter.

-key *prepare | commit | cancel*

(Optional) Manages the creation of a new or replacement encryption keys when **-usb newkey** is specified.

-key *prepare*

Generates system encryption keys and writes those keys to all system attached USB flash drives. If there is active encryption key material, confirm that at least one USB flash drive has the current key material. Use this command only when the **lsencryption** value for **usb_rekey** is set to **no** or **no_key**.

-key *commit*

Commits the prepared key as the current key. Use this command when the **lsencryption** value for **usb_rekey** is set to **prepared** and the number of USB encryption keys is at least the minimum number required.

-key *cancel*

Cancels any specified key changes. Use this command when the **lsencryption** value for **usb_rekey** is set to **prepared**.

Description

Use this command to manage the encryption state of the system.

It can turn on or turn off USB key encryption or key server encryption (but you cannot disable encryption if there are any encrypted arrays). There are four types:

- **enable**, which enables encryption
- **disable**, which disables encryption
- **validate**, which validates USB key encryption

Note: The **validate** option does not apply to key server authentication.

- **newkey**, which specifies a new USB key for encryption

You can also rekey the external USB key or key server key information, which is divided into three stages:

- **prepare**, which generates new keys and sets up the system to change encryption keys during **apply**
- **commit**, which includes applying new keys (and copying key information)

- `cancel`, which rolls back the key setup performed during the `prepare` and cancels the `rekey` request

Specify `chencryption -keyserver newkey -key commit` to add a new key on a key server.

Note: You cannot use USB flash drive authentication and key server authentication in parallel on the same system.

An invocation example

```
chencryption -usb enable
```

The resulting output:

No feedback

An invocation example

```
chencryption -usb newkey -key prepare
```

The resulting output:

No feedback

An invocation example

```
chencryption -usb newkey -key commit
```

The resulting output:

No feedback

An invocation example

```
chencryption -keyserver enable
```

The resulting output:

```
chencryption -keyserver newkey -key prepare
```

An invocation example

```
chencryption -keyserver newkey -key commit
```

The resulting output:

No feedback

An invocation example

The resulting output:

No feedback

lsencryption

Use the **lsencryption** command to display system encryption information.

Syntax

```

>> lsencryption — [ -nohdr ] [ -delim delimiter ]

```

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) In a detailed view, each item of data has its own row, and if the headings are displayed, the data is separated from the heading by a space. The **-delim** parameter overrides this behavior. Valid input for the **-delim** parameter is a one-byte character. In a detailed view, the data is separated from its heading by the specified delimiter.

Description

Use this command to display output related to the system encryption state.

Table 76 describes possible outputs.

Table 76. Isencryption output

Attribute	Value
status	Indicates the system encryption status. <ul style="list-style-type: none">• not_supported, which indicates that the system has no supported encryption function.• not_licensed, which indicates that the system supports encryption but not all licenses are installed.• licensed, which indicates that the system has licenses installed for all encryption-capable hardware.• enabled , which indicates that system encryption is working and ready to create encrypted storage.
error_sequence_number	Indicates the event log sequence number of any problem affecting encryption. If there is no problem it is blank.
usb_rekey	Indicates the state of the Universal Serial Bus (USB) rekey process. <ul style="list-style-type: none">• no, which indicates that there is no rekey process ongoing, but keys exist.• no_key, which indicates that there is no rekey process and keys do not exist.• prepared, which indicates that a rekey process is active and the system has prepared a new key that is waiting for this command to be issued: chencryption -usb newkey -key commit.• committing, which indicates that a commit is in progress.
usb_key_copies	Indicates the number of USB devices that prepared keys have been written to. The value must be a numeric string.
usb_key_filename	Indicates the name of the file containing the current encryption key. The value must be an alphanumeric string containing between 1 and 110 ASCII characters.
usb_rekey_filename	Indicates the name of the file containing the current prepared encryption key.
keyserver_status	Indicates the encryption status for key server encryption. The values are: <ul style="list-style-type: none">• not_supported, which indicates that the system has no supported encryption function.• not_licensed, which indicates that the system supports encryption but not all licenses are installed.• licensed, which indicates that the system has licenses installed for all encryption-capable hardware.• enabled , which indicates that system encryption is working and ready to create encrypted storage.

Table 76. *lsencryption output (continued)*

Attribute	Value
keyserver_rekey	Indicates the state of the key server rekey process. The values are: <ul style="list-style-type: none"> no, which indicates that there is no rekey process ongoing, but keys exist. no_key, which indicates that there is no rekey process and keys do not exist. prepared, which indicates that a rekey process is active and the system has prepared a new key that is waiting for this command to be issued: chencryption -keyserver newkey -key commit. committing, which indicates that a commit is in progress.
keyserver_pmk_uid	Indicates the UID for the key server.
keyserver_rekey_pmk_uid	Indicates the UID (after a rekey process) for the key server.

An invocation example for an encrypted system with no rekey

```
lsencryption
```

The resulting output:

```
status enabled
error_sequence_number
usb_rekey no
usb_key_copies 0
usb_key_filename
usb_rekey_filename
keyserver_status disabled
keyserver_rekey no_key
keyserver_pmk_uid
keyserver_rekey_pmk_uid
```

An invocation example for an encrypted system during the rekey

```
lsencryption
```

The resulting output:

```
status enabled
error_sequence_number
usb_rekey prepared
usb_key_copies 3
usb_key_filename
usb_rekey_filename encryptionkey_0000020061800028_0010030C00000007_Cluster_9.19.88.231
keyserver_status enabled
keyserver_rekey prepared
keyserver_pmk_uid
keyserver_rekey_pmk_uid KEY-1b9dcbe7-8b1c-401d-9bc2-1791534689fc
```

An invocation example for an encrypted system after the rekey completes

```
lsencryption
```

The resulting output:

```
status enabled
error_sequence_number
usb_rekey no
usb_key_copies 3
usb_key_filename encryptionkey_0000020061800028_0010030C00000007_Cluster_9.19.88.231
usb_rekey_filename
```


Chapter 14. Licensing and featurization commands

The licensing commands enable you to work with SAN Volume Controller licensed functions.

activatefeature

Use the **activatefeature** command to activate a feature (using a license key or keyfile) or feature trial period.

Syntax



Parameters

-trial *feature_id* | *feature_name*

(Optional) Activates the trial period for the feature of the specified ID that uses an unsigned 16-bit integer:

- Valid integer values are 0, 1, and 3.
- Valid names are `turbo_performance`, `easy_tier`, and `remote_mirroring`.

-licensekey *key*

(Optional) Provides the license key to activate a feature that contains 16 hexadecimal characters organized in four groups of four numbers with each group separated by a hyphen (such as 0123-4567-89AB-CDEF).

-licensekeyfile *filepath*

(Optional) Provides the full path-to-file containing all required license information by using an alphanumeric string that contains 1 - 256 characters.

Description

All parameters are mutually exclusive.

A license key file can contain one or more license keys. If you specify a key file, every key in the file is applied to the system. The license key is checked against the node or control enclosure serial number, machine type, and model. If no valid keys exist in the file, the command cannot complete successfully on the system. If you cannot apply a key successfully to the system, the command adds any remaining keys.

You must have one key for each node or control enclosure. Specify `activatefeature -licensekeyfile` with an `.xml` file that contains all node or control enclosure keys. Or, specify `activatefeature -licensekey` one time per node or control enclosure.

If a feature is already activated and you activate a feature again by using a key, the command completes successfully.

Remember:

- You cannot complete a trial when a feature is activated.

- You can activate a feature while a trial is in progress.

An invocation example

```
activatefeature -trial 1
```

The resulting output:

Activation of a trial is a one time operation. Are you sure you wish to continue? Yes

An invocation example

```
activatefeature -licensekey 0123-4567-89AB-CDEF
```

The resulting output:

No feedback

An invocation example

```
activatefeature -licensekeyfile /tmp/keyfile.xml
```

The resulting output:

No feedback

chlicense

Use the **chlicense** command to change license settings for clustered system (system) features.

Syntax

This applies only to SAN Volume Controller:

```

>> chlicense — [ -flash — capacity_TB ————— ] [ -remote — capacity_TB ————— ] [ -cloud ]
                  | [ -virtualization — virtualization_setting_in_scu ————— ]
                  | [ -compression — compression_setting_in_scu ————— ]
                  | [ -physical_flash — [ on ————— ]
                  |                               [ off ————— ]

```

This applies only to Storwize V7000:

```

>> chlicense — [ -remote ————— ]
                  | [ -virtualization ————— ]

```

This applies only to Storwize V5000.

```

>> chlicense — [ -remote — enclosures ————— ]
                  | [ -virtualization — number_of_enclosures ————— ]
                  | [ -easytier — enclosures ————— ]
                  | [ -flash — enclosures ————— ]

```

Parameters

This applies to SAN Volume Controller, Storwize V7000:

-flash *capacity_TB*

(Optional) Changes system 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 system licensing for remote copy features such as Metro Mirror, Global Mirror, and HyperSwap. To change the licensed capacity for one of these features, specify a capacity in terabytes (TB). You must have an enclosure license for all enclosures.

Note: For Storwize V7000, Flex V7000 Storage Node, Storwize V3500, and Storwize V3700, specify the total number of internal and external enclosures that you have licensed on your system.

-virtualization *virtualization_setting_in_scu*

(Optional) Changes system 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 to use.

-physical_flash **on | off**

(Optional) For physical disk licensing, enables or disables the FlashCopy feature. The default value is **off**.

-compression *virtualization_setting_in_scu*

(Optional) Changes system licensing for the compression feature.

Note: All Storwize V7000 systems support compression.

To change the compression license capacity, specify a capacity value in terabytes (TB).

-cloud *enclosures*

(Optional) Specifies number of enclosures for the transparent cloud tiering feature. The value must be a number.

This applies only to Storwize V5000.

-remote *enclosures*

(Optional) Changes system licensing for the Metro Mirror, Global Mirror, or HyperSwap feature, and specifies the total number of internal and external enclosures that you have licensed on your system. You must have a Metro and Global Mirror enclosure license for all enclosures.

-virtualization *enclosures*

(Optional) Changes system licensing for the Virtualization feature. Specify the number of enclosures of external storage that you have been authorized to use.

-easytier *enclosures*

(Optional) Specifies the enclosures on which the customer can run Easy Tier.

-flash *enclosures*

(Optional) Specifies the total number of internal and external enclosures for the FlashCopy feature.

-compression *enclosures*

(Optional) Changes system licensing for the compression feature.

Description

The **chlicense** command changes license settings for the system. Any change that is made is logged as an event in the license setting log.

The capacity for each licensed feature can be modified with this command. This is the number of terabytes (TB) of volume capacity or Storage Control Units (SCU) capacity that can be configured by the system.

The enclosure license already includes virtualization of internal drives on your system. You can use this command to set any additional options. The total number of enclosures in your system must not exceed the total number of licensed enclosures that you have. The total virtualized capacity (number of external enclosures that can be configured by the system) can also be modified with this command. The default is to have no feature licensed, but this does not stop you from using related functions.

Any error that is placed in the license settings log results in a generic error being placed in the system error log. The command-line tool return code also notifies you that you are using an unlicensed feature.

When you reach 90% capacity, any attempt to create or extend volumes, relationships, or mappings generates an error message. You can still create and expand volumes, relationships, or mappings. When usage reaches or exceeds 100% capacity, errors are placed in the license settings log stating that you are using an unlicensed feature.

An invocation example for adding a remote copy license capacity of 5 TB

```
chlicense -remote 5
```

The resulting output:

No feedback

An invocation example for enabling Easy Tier settings

```
chlicense -easytier 2
```

The resulting output:

No feedback

An invocation example for modifying a compression license value

```
chlicense -compression 4
```

The resulting output:

No feedback

An invocation example for changing the license on a cloud account

```
chlicense -cloud 2
```

The resulting output:

No feedback

deactivatefeature

Use the **deactivatefeature** command to deactivate a feature or suspend a feature trial period.

Syntax

```
➤ deactivatefeature — feature_id — ➤
```


Parameters

feature_id

(Required) Deactivates the feature (or feature trial). This is the unique ID as displayed when using the **lsfeature** command, and is an incremental number (from 0 to 320).

Description

Use this command to deactivate a feature or suspend a feature trial period.

An invocation example

```
deactivatefeature 1
```

The resulting output:

```
You are removing the ability to use a feature of this system. Are you sure you wish to continue? Y
```

lsfeature

Use the **lsfeature** command to list the features that are available for the current clustered system (system) code release. You can also list trial or entitlement information and license keys.

Syntax

```
►► lsfeature [ -delim delimiter ] [ -nohdr ] [ -bytes ] ►►
```

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.

-bytes

(Optional) Displays all capacities as bytes.

Description

This command lists the features that are available for the current clustered system (system) code release. You can also list trial or entitlement information and license keys.

This table provides the attribute values that can be displayed as output view data.

Table 77. *lsfeature* outputs

Attribute	Possible Values
id	Indicates the unique ID (2-character) feature number.
name	Indicates the feature name with a 16-character alphanumeric string: <ul style="list-style-type: none"> • easy_tier • remote_mirroring • flashcopy_upgrade • turbo_performance • encryption
state	Indicates the current state of the feature: <ul style="list-style-type: none"> • active • inactive • trial_available • trial_active • trial_expired
license_key	Indicates the key used in feature activation with a string consisting of 16 hexadecimal characters organized in four groups of four numbers with each group separated by a hyphen (such as 0123-4567-89AB-CDEF).
trial_expiration_date	Indicates the trial expiration date as long as the state is trial_available or trial_active. This value must be in the format YYYYMMDD.
serial_num	Indicates the product serial number.
mtm	Indicates the machine type and model.

Note: A license key's association with an enclosure can be determined if:

- The enclosure associated with the key contains at least one node that has been added to the cluster and the node is either online or offline at the time the view is queried.
- The enclosure associated with the key contains at least one node that is currently a candidate for clustered system membership. The node must be online at the time the view is being queried.

An invocation example

`lsfeature`

The resulting output:

id	name	state	license_key	trial_expiration_date	serial_num	mtm
0	turbo_performance	trial_available		20130201		
1	easy_tier	trial_active		20130101		
2	flashcopy_upgrade	active	0123-4567-89AB-CDEF			
3	remote_mirroring	trial_expired		20130201		

An invocation example

In this system, both licenses are for encryption. There are 2 control enclosures, and the serial number and machine type are displayed:

`lsfeature`

The resulting output:

id	name	state	license_key	trial_expiration_date	serial_num	mtm
0	encryption	active	90AB-D41D-C799-2EF4		78G00TT	2076-112
1	encryption	active	3A87-463E-B5DF-9969		31G00KG	2076-124

If one of the licenses is removed:

lsfeature

The resulting output:

id	name	state	license_key	trial_expiration_date	serial_num	mtm
0	encryption	inactive	90AB-D41D-C799-2EF4		78G00TT	2076-112

The state is inactive because control enclosures require their own license to activate encryption.

An invocation example

In this system, both licenses are for encryption. There are 2 control enclosures, and the serial number and machine type are displayed:

lsfeature

The resulting output:

id	name	state	license_key	trial_expiration_date	serial_num	mtm
0	encryption	active	90AB-D41D-C799-2EF4		78G00TT	2076-112
1	encryption	active	3A87-463E-B5DF-9969		31G00KG	2076-124

If one of the licenses is removed:

lsfeature

The resulting output:

id	name	state	license_key	trial_expiration_date	serial_num	mtm
0	encryption	inactive	90AB-D41D-C799-2EF4		78G00TT	2076-112

The state is inactive because control enclosures require their own license to activate encryption.

Islicense

Use the **lslicense** command to display current license settings for clustered system (system) features.

Syntax

```

>> lslicense — [ -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 system 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, Flex System V7000 Storage Node, Storwize V3500, and Storwize V3700 lists enclosure license values.

Use the **chlicense** command to change the feature license settings. Because the feature license settings are entered when the system is first created, you must only update the settings if you have changed your license.

Table 78 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 78. lslicense output

Attribute	Possible Values
used_flash	Indicates the amount Flash Copy (FC) memory used.
used_remote	Indicates the amount of remote copy memory used.
used_virtualization	Indicates the amount of virtualization memory used.
license_flash	Indicates the FC license settings.
license_remote	Indicates remote copy license settings.
license_virtualization	Indicates license virtualization settings.
license_physical_disks	Indicates the amount of physical disk space available for the license.
license_physical_flash	Indicates if the license physical flash is on or off.
license_physical_remote	Indicates if the license physical remote is on or off.
used_compression_capacity	Indicates the total virtual size of volumes with compressed copies, in total bytes (numeric format with two decimal places).
license_compression_capacity	Indicates the licensed compression capacity, in total bytes (numeric format).
license_compression_enclosures	Indicates which licensed enclosures have compression (numeric format).
license_easy_tier	Indicates which enclosures Easy Tier can be run on.
license_cloud_enclosures	Indicates whether a separate cloud account system storage license is configured.
scu_ratio_ssd	Indicates the storage capacity unit (SCU) ratio for SSD tier storage. The value must be a number with 2 decimal places.
scu_ratio_enterprise	Indicates the SCU ratio for enterprise tier storage. The value must be a number with 2 decimal places.
scu_ratio_nearline	Indicates the SCU ration for nearline tier storage. The value must be a number with 2 decimal places.

An invocation example

```
lslicense
```

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 on  
license_physical_remote off  
used_compression_capacity 0.02  
license_compression_capacity 0  
license_compression_enclosures 1
```

```
license_cloud_enclosures 0  
scu_ratio_ssd 1.00  
scu_ratio_enterprise 1.18  
scu_ratio_nearline 4.00
```


Chapter 15. FlashCopy commands

The FlashCopy commands enable you to work with FlashCopy methods and functions with the SAN Volume Controller.

chfcconsistgrp

Use the **chfcconsistgrp** command to change the name of a consistency group or marks the group for auto-deletion.

Syntax



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 **chfcconsistgrp** command changes the name of a consistency group, marks the group for auto-deletion, or both.

Note: Maps that are *rc_controlled* are not shown in the view when this command is specified.

An invocation example

```
chfcconsistgrp -name testgrp1 fcconsistgrp1
```


The resulting output:


No feedback

chfcmap

Use the **chfcmap** command to modify attributes of an existing mapping.

Syntax





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.

Important: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information.

-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.

-copyrate *rate*

(Optional) Specifies the copy rate. The *rate* value can be 0 - 150. 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 79 on page 431.

-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 - 150. 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 **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 volumes 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 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.

This table 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 79. Relationship between the cleaning rate, data rate, and grains per second values

User-specified cleaning <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
101 - 110	128 MB	512	2048
111 - 120	256 MB	1024	4096
121 - 130	512 MB	2048	8192
131 - 140	1 GB	4096	16384
141 - 150	2 GB	8192	32768

Note: Maps that are `rc_controlled` are not shown in the view when this command is specified.

An invocation example

```
chfcmap -name testmap 1
```

The resulting output:

No feedback


lsfcconsistgrp

Use the **lsfcconsistgrp** command to display a concise list or a detailed view of FlashCopy consistency groups that are visible to the clustered system (system). 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 system. (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:

```
lsfcconsistgrp -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 **lsfcconsistgrp** 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 system.

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

id Displays the mapping ID.

name Displays the mapping name.

start_time
Specifies the time the group was started in YYMMDDHHMMSS format (or blank).

autodelete
Specifies if autodelete is on or off.

FC_mapping_id
Displays the FlashCopy mapping ID.

FC_mapping_name
Displays the FlashCopy mapping name.

estimated_copying_completion_time
Displays the estimated time that the background copy takes to complete. It is in the format YYMMDDHHMMSS. If a copy has not been started or the copyrate is 0, a blank line is displayed

estimated_stopping_completion_time
Displays the estimated time that the stopping takes to complete. It is in the format YYMMDDHHMMSS. If a copy is not stopping a blank line is displayed.

A concise invocation example

```
lsfcconsistgrp -delim :
```

The concise resulting output:

```
id:name:status:start_time
1:ffccg0:empty:060627083237
2:ffccg1:idle_or_copied:060627083337
3:ffccg2:idle_or_copied:060627083437
```

A detailed invocation example

```
lsfcconsistgrp -delim : 1
```

The detailed resulting output:

```
id:1
name:ffccg0
status:empty
```

A detailed invocation example

```
lsfcconsistgrp -delim : fccstgrp0
```

The detailed resulting output:


```
id:1
name:FCcgrp0
status:idle_or_copied
start_time:060627083137
autodelete:off
FC_mapping_id:0
FC_mapping_name:fcmap0
```

```
FC_mapping_id:1
FC_mapping_name:fcmap1
esimated_copying_completion_time:
esimated_stopping_completion_time:
```

lsfcmap

Use the **lsfcmap** command generate a list containing concise information about all of the FlashCopy mappings that are visible to the clustered system (system), or detailed information for a single FlashCopy mapping.

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 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 (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:

```
lsfcmap -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 system.

The following list shows attribute values that can be displayed as output view data:

- id** Displays the mapping ID.
- name** Displays the mapping name.
- source_vdisk_id**
Displays the source volume ID.
- source_vdisk_name**
Displays the source volume name.
- target_vdisk_id**
Displays the target volume ID.
- target_vdisk_name**
Displays the target volume name.
- group_id**
Displays the group ID.
- group_name**
Displays the group name.
- status** Displays the status:
 - idle_or_copied
 - preparing
 - prepared
 - copying
 - stopped
 - suspended
 - stopping

progress
Displays the progress.

copy_rate
Displays the copy rate.

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.

dependent_mappings
Displays any dependent mappings.

autodelete
Specifies if autodelete is on or off.

clean_progress
Indicates the clean progress.

clean_rate
Indicates the clean rate.

incremental
Indicates whether incremental is on or off.

difference
Indicates the difference.

IO_group
Displays the I/O group ID.

IO_group_name
Displays the I/O group name.

partner_FC_id
Displays the partner FlashCopy ID/

partner_FC_name
Displays the partner FlashCopy name.

restoring
Displays the restoring status. The values are yes or no.

rc_controlled
Displays the rc_controlled status.

copy_rate_mb
Displays the copy rate MB amount.

clean_rate_mb
Displays the clean rate MB amount.

keep_target
Displays the target and source volume availability. The values are:

- yes, which indicates that the source volume availability is connected to the availability of the target volume
- no, which indicates that if there is a problem with the target volume that can impact FlashCopy operations, the target volume is removed

estimated_copying_completion_time
Displays the estimated time that the background copy takes to complete. It is in the format YYMMDDHHMMSS. If a copy has not been started or the copyrate is 0, a blank line is displayed

estimated_stopping_completion_time

Displays the estimated time that the stopping takes to complete. It is in the format YYMMDDHHMMSS. If a copy is not stopping a blank line is displayed.

restore_progress

Displays the percentage of the source volume that has been restored from the target.

estimated_restore_completion_time

Displays the estimated time that the restore takes to complete. It is in the format YYMMDDHHMMSS. If a copy is not restoring then a blank line is displayed.

Note: Using `rc_controlled` indicates that the map is for internal use only. It cannot be manipulated externally.

A concise invocation example

```
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:start_time:rc_controlled
0:test:0:vdisk0:1:vdisk1:idle_or_copied:0:50:100:off:no
no0: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

```
lsfcmap 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
rc_controlled:no
keep_target:yes
esimated_copying_completion_time:
esimated_stopping_completion_time:
restore_progress:
esimated_retores_completion_time:
```

lsfcmapcandidate

Use the **lsfcmapcandidate** command to list all of the volumes that are associated with fewer than 256 FlashCopy mappings.

Syntax

```
►► — 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 volumes that are associated with fewer than 256 FlashCopy mappings.

An invocation example

```
lsfcmapcandidate
```

The resulting output:

```
id
2
3
4
```

lsfcmapprogress

Use the **lsfcmapprogress** command to display the progress of the background copy of a FlashCopy mapping. This information is displayed as a percentage-completed value.

Syntax

```
►► — lsfcmapprogress — [ -nohdr ] [ -delim — delimiter ] [ fcmap_id ] [ fcmap_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

```
lsfcmapprogress 0
```

The resulting output:

id	progress
0	0

lsfcmapdependentmaps

Use the **lsfcmapdependentmaps** command to display the FlashCopy mappings that are dependent on the user specified mapping.

Syntax

```
➤— lsfcmapdependentmaps — [ -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.

An invocation example

```
lsfcmapdependentmaps -delim : 2
```

The resulting output:

```
fc_id:fc_name
1:fcmap1
3:fcmap3
```

lsrmvdiskdependentmaps

Use the **lsrmvdiskdependentmaps** command to display all FlashCopy mappings that must be stopped for the specified volume to be deleted.

Syntax

```
lsrmvdiskdependentmaps [-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 volume for which the FlashCopy mappings are displayed.

Description

This command returns a list of the FlashCopy mappings that must be stopped before the specified volume can be deleted. Any mappings that are returned in the list for the volume are automatically

stopped when the volume is deleted with the **force** option.

An invocation example

```
lsrmvdiskdependentmaps -delim : 0
```

The resulting output:

```
id:name
2:fcmap2
5:fcmap5
```

mkfcconsistgrp

Use the **mkfcconsistgrp** command to create a new FlashCopy consistency group and identification name.

Syntax

```
➤— 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`.

Note: Consistency group names must be an alphanumeric string of up to 15 characters.

-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 volumes 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.

Note: Maps that are `rc_controlled` are not shown in the view when this command is specified.

Remember: Names representing Metro Mirror or Global Mirror consistency groups relationships are restricted to fifteen characters in length (not sixty-three for an extended character set).

An invocation example

```
mkfcconsistgrp
```

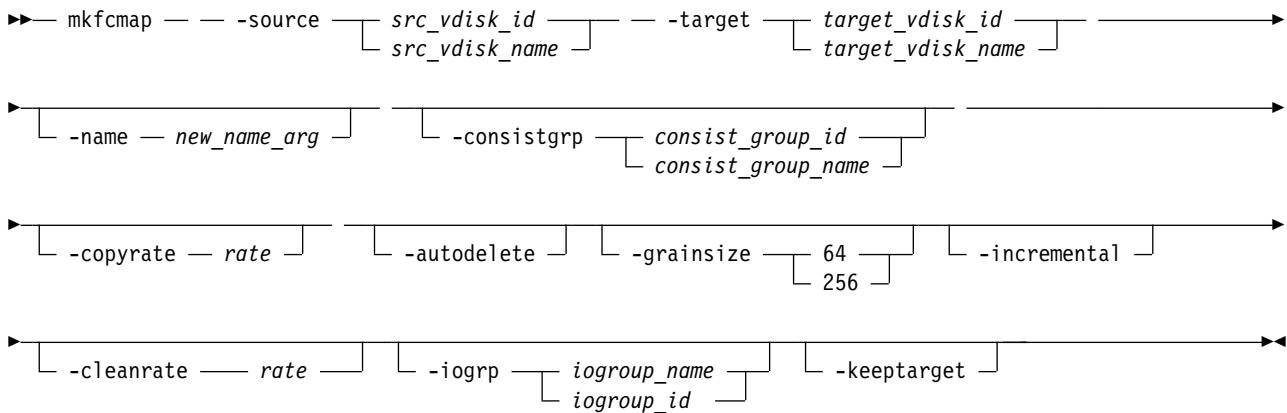
The resulting output:

```
FlashCopy Consistency Group, id [1], successfully created
```

mkfcmap

Use the **mkfcmap** command to create a new FlashCopy mapping, which maps a source volume to a target volume for subsequent copying.

Syntax



Parameters

- source** *src_vdisk_id* | *src_vdisk_name*
(Required) Specifies the ID or name of the source volume.
- target** *target_vdisk_id* | *target_vdisk_name*
(Required) Specifies the ID or name of the target 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 - 150. 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 80 on page 444.
- 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.

Remember: If either the source or target disk contains compressed copies, the default value is 64 (unless source or target disk is part of a mapping with grainsize 256 KB).
- 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 - 150. 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 volume, if a single target map, or the I/O group of the other FlashCopy mapping to which either the source or target 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).

-keeptarget

(Optional) Specifies that the target volume and source volume availability should be kept the same. If the target becomes unavailable the source is also made unavailable (instead of stopping the FlashCopy mapping).

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 is set to on. The source and target volumes must be specified on the **mkfcmap** command. The **mkfcmap** command fails if the source and target volumes are not identical in size. Issue the **lsvdisk -bytes** command to find the exact size of the source volume for which you want to create a target disk of the same size. The target volume that you specify cannot be a target 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 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 volumes are the target and source 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.

You can create a FlashCopy mapping in which the target volume is a member of a Metro Mirror or Global Mirror relationship, unless one of the following conditions applies:

- The relationship is with a clustered system that is running an earlier code level.
- The I/O group for the mapping is different than the I/O group for the proposed mapping target volume.

Note: You cannot use this command if a volume is part of a mapping and cloud snapshot is enabled on the volume.

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 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.

Note: You cannot issue this command if the:

- Target volume is a master or auxiliary volume in an active-active relationship

- Source volume is a master or auxiliary volume in an active-active relationship and the target volume and map are not in the same site as the source volume

This table 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 80. Relationship between the cleaning rate, data rate, and grains per second values

User-specified cleaning rate 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
101 - 110	128 MB	512	2048
111 - 120	256 MB	1024	4096
121 - 130	512 MB	2048	8192
131 - 140	1 GB	4096	16384
141 - 150	2 GB	8192	32768

Note: Maps that are rc_controlled are not shown in the view when this command is specified.

An invocation example

```
mkfcmap -source 0 -target 2 -name mapone
```

The resulting output:

```
FlashCopy Mapping, id [1], successfully created
```

An invocation example

```
mkfcmap -source 0 -target 2 -name mapone -keeptarget
```

The resulting output:

```
FlashCopy Mapping, id [1], successfully created
```

prestartfcconsistgrp

Use the **prestartfcconsistgrp** command to prepare 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 volume and forces the cache into the write-through mode until the consistency group is started.

Syntax

```

▶▶ 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 volume of one of the mappings in the consistency group is being used as a source volume 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 volume is first flushed to disk. This step ensures that the FlashCopy target volume is identical to what has been acknowledged to the host operating system as having been written successfully to the source volume.

You can use the **restore** parameter to force the consistency group to be prepared even if the target volume of one or more mappings in the consistency group is being used as a source volume of another active mapping. In this case the mapping restores as shown in the **lsfcmap** view. If the **restore** parameter is specified when preparing a consistency group where none of the target volumes are the source volume of another active mapping, then the parameter is ignored.

You must issue the **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 consistency group must be prepared before it can be started.

This command is rejected if the target of a FlashCopy mapping in the consistency group is in a remote copy relationship, unless the relationship is one of the following types and is the secondary target of the remote copy:

- **idling**
- **disconnected**
- **consistent_stopped**
- **inconsistent_stopped**

The FlashCopy mapping also fails in the following cases:

- You use the **prep** parameter.
- The target volume is an active remote copy primary or secondary volume.

- The FlashCopy target (and remote copy primary target) volume is offline. If this occurs, the FlashCopy mapping stops and the target volume remains offline.

Note: Maps that are `rc_controlled` are not shown in the view when this command is specified.

An invocation example

```
prestartfcconsistgrp 1
```

The resulting output:

No feedback

prestartfcmap

Use the **prestartfcmap** command to prepare a FlashCopy mapping so that it can be started. This command flushes the cache of any data that is destined for the source volume and forces the cache into the write-through mode until the mapping is started.

Syntax

```

➔ 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 volume is being used as a source volume 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 volume 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 volume is being used as a source volume of another active mapping. In this case, the mapping is restoring as shown in the **lsfcmap** view. If the **restore** parameter is specified when preparing a mapping where the target volume is not the source volume 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. For example, while a volume is in the prepared state, response times might increase.

This command is rejected if the target of the FlashCopy mappings is the secondary volume in a Metro Mirror, Global Mirror, or active-active relationship (so that the FlashCopy target is the remote copy secondary). Remote copy includes Metro Mirror, Global Mirror, and active-active.

Note: If the remote copy is idling or disconnected, even if the FlashCopy and remote copy are pointing to the same volume, the auxiliary volume is not necessarily the secondary volume. In this case, you can start a FlashCopy mapping.

The FlashCopy mapping also fails in the following cases:

- The remote copy is active.
- The FlashCopy target (and remote copy primary target) volume is offline. If this occurs, the FlashCopy mapping stops and the target volume remains offline.

Note: Maps that are `rc_controlled` are not shown in the view when this command is specified.

An invocation example

```
prestartfcmap 1
```

The resulting output:

No feedback

rmfcconsistgrp

Use the **rmfcconsistgrp** command to delete a FlashCopy consistency group.

Syntax

```
➤— 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.

Important: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information.

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 **rmfcmap** command.

Note: Maps that are `rc_controlled` are not shown in the view when this command is specified.

An invocation example

```
rmfcconsistgrp fcconsistgrp1
```

The resulting output:

No feedback

rmfcmap

Use the **rmfcmap** command to delete an existing mapping.

Syntax

```
➤— rmfcmap — [ -force ] [ fc_map_id | fc_map_name ] —➤
```

Parameters

-force

(Optional) Specifies that the target volume 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 volumes; it does not affect the volumes themselves. However, if you force the deletion, the target volume (which might contain inconsistent data) is brought back online.

If the target of the FlashCopy mapping is a member of the remote copy created, the remote copy can be affected in the following ways:

- If a stopped FlashCopy mapping is deleted and the I/O group associated with the FlashCopy mapping is suspended while this delete is being processed, then all remote copy relationships associated with the target volume of a the FlashCopy mapping that were active while the FlashCopy mapping was copying can be corrupted. You must resynchronize them next time you start the system.
- If a stopped FlashCopy mapping that has previously failed to prepare is deleted, then all remote copy relationships in the set of remote copy relationships associated with the target volume can be corrupted. You must resynchronize them next time you start the system.

Note: Remote copy includes Metro Mirror, Global Mirror, and HyperSwap.

Note: Maps that are `rc_controlled` are not shown in the view when this command is specified.

An invocation example

```
rmfcmap testmap
```

The resulting output:

No feedback

startfcconsistgrp

Use the **startfcconsistgrp** command to start a FlashCopy consistency group of mappings. This command makes a point-in-time copy of the source volumes at the moment that the command is started.

Syntax

►► 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 volume of one of the mappings in the consistency group is being used as a source volume 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 volumes 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 volume of one or more mappings in the consistency group is being used as a source volume of another active mapping. In this case, the mapping is restoring as shown in the **lsfcmap** view. If the **restore** parameter is specified when starting a consistency group where none of the target volumes are the source volume of another active mapping, the parameter is ignored.

If a consistency group is started and the target volume 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 volume, 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 issue the **startfcconsistgrp** command to start the copy.

This command is rejected if the target of the FlashCopy mapping in the specified consistency group is the secondary volume in a remote copy relationship (so that the FlashCopy target is the remote copy secondary).

Note: If the remote copy is idling or disconnected, even if the FlashCopy and remote copy are pointing to the same volume, the auxiliary volume is not necessarily the secondary volume. In this case, you can start a FlashCopy mapping.

The FlashCopy mapping also fails in the following cases, if the target of the FlashCopy mapping in the specified consistency group is the primary volume in a remote copy relationship (so that the FlashCopy target is the remote copy primary target):

- The remote copy is active.
- The FlashCopy target (and remote copy primary target) volume is offline. If this occurs, the FlashCopy mapping stops and the target volume remains offline.

Maps that are `rc_controlled` are not shown in the view when this command is specified.

If any source volumes in the FlashCopy consistency group are in an active-active relationship, the group can only be started if the information on all those source volumes current, or an older copy to which access has been provided by specifying:

```
stoprcrelationship -access
```

A current volume in an active-active relationship is the primary copy, or the secondary copy when the relationship's state is `consistent_synchronized`.

An invocation example

```
startfcconsistgrp -prep 2
```


The resulting output:

No feedback

startfcmap

Use the **startfcmap** command to start a FlashCopy mapping. This command makes a point-in-time copy of the source volume at the moment that the command is started.

Syntax



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.

Note: If you have already used the **prestartfcmap** command, you cannot use the **-prep** parameter on the **startfcmap** command; the command fails. However, if the FlashCopy has successfully prepared before, the **startfcmap** command succeeds.

-restore

(Optional) Specifies the restore flag. When combined with the **prep** option, this forces the mapping to be prepared even if the target volume is being used as a source volume in another active mapping. An active mapping is in the copying, suspended, or stopping state.

fc_map_id | *fc_map_name*

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 volume. You can combine the **restore** parameter with the **prep** parameter to force the mapping to be prepared prior to starting, even if the target volume is being used as a source volume of another active mapping. In this case, the mapping is restoring as shown in the **lsfcmap** view. If the **restore** parameter is specified when starting a mapping where the target volume is not the source volume of another active mapping, the parameter is ignored and the mapping is not restoring as shown in the **lsfcmap** view.

If a mapping is started and the target volume 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 volume, 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.

This command is rejected if the target of the FlashCopy mapping is the secondary volume in a Metro Mirror or Global Mirror relationship (so that the FlashCopy target is the remote copy secondary).

Note: If the remote copy is idling or disconnected, even if the FlashCopy and remote copy are pointing to the same volume, the auxiliary volume is not necessarily the secondary volume. In this case, you can start a FlashCopy mapping.

The FlashCopy mapping also fails in the following cases, if the target of the FlashCopy mapping is the primary volume in a Metro Mirror or Global Mirror relationship (so that the FlashCopy target is the remote copy primary):

- The remote copy is active.
- The FlashCopy target (and remote copy primary target) volume is offline. If this occurs, the FlashCopy mapping stops and the target volume remains offline.

Maps that are `rc_controlled` are not shown in the view when this command is specified.

Remember: If the source volume is in an active-active relationship then the FlashCopy mapping can only be started if the information on the source volume is current, or an older copy to which access has been provided by specifying:

```
stoprcrelationship -access
```

A current volume in an active-active relationship is the primary copy, or the secondary copy when the relationship's state is `consistent_synchronized`.

An invocation example

```
startfcmap -prep 2
```

The resulting output:

No feedback

stopfcconsistgrp

Use the **stopfcconsistgrp** command to stop 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

```
stopfcconsistgrp [-force|-split] fc_consist_group_id_or_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.

If the **-force** parameter is not specified, the command is rejected if the target volume of the FlashCopy consistency group is the primary in a relationship that is mirroring I/O:

- consistent_synchronized
- consistent_copying
- inconsistent_copying

If the **-force** parameter is specified, any Metro Mirror or Global Mirror relationships associated with the target volumes of the FlashCopy mappings in the specified consistency group stops. If a remote copy relationship associated with the target was mirroring I/O when the map was copying, it might lose its difference recording capability and require a full resynchronization upon a subsequent restart.

-split

(Optional) Breaks the dependency on the source volumes of any mappings that are also dependent on the target volume. 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 **lsfcmap** command.

fc_consist_group_id_or_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 volume is reported as offline if it does not contain a complete image. Before you can access this volume, the group of mappings must be prepared and restarted.

If the consistency group is in the *idle_or_copied* state, the **stopfcconsistgrp** command has no effect and the consistency group stays in the *idle_or_copied* state.

Note: Prior to SAN Volume Controller 4.2.0, the **stopfcconsistgrp** command always caused the consistency group to go to the stopped state, taking the target volumes 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 volumes. 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.

Note: Maps that are *rc_controlled* are not shown in the view when this command is specified.

An invocation example

```
stopfcconsistgrp testmapone
```

The resulting output

No feedback

stopfcmap

Use the **stopfcmap** command to stop all processing that is associated with a FlashCopy mapping that is in one of the following processing states: prepared, copying, stopping, or suspended.

Syntax

```
➤— stopfcmap — [ -force -split ] — fc_map_id_or_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.

If the **-force** parameter is not specified, the command is rejected if the target volume of the FlashCopy mapping is the primary in a relationship which is mirroring I/O:

- consistent_synchronized
- consistent_copying
- inconsistent_copying

If the **-force** parameter is specified to a FlashCopy mapping whose target volume is also in a Metro Mirror or Global Mirror relationship, the relationship stops. If a remote copy relationship associated with the target was mirroring I/O when the map was copying, it might lose its difference recording capability and require a full resynchronization on a subsequent restart.

-split

(Optional) Breaks the dependency on the source volume 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 **lsfcmap** command.

fc_map_id_or_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 volume 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 volume. 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.

Note: Maps that are `rc_controlled` are not shown in the view when this command is specified.

Remember: If the source volume is in an active-active relationship then the FlashCopy mapping can only be stopped if the information on the source volume is current, or an older copy to which access has been provided by specifying:

```
stopprrelationship -access
```

A current volume in an active-active relationship is the primary copy, or the secondary copy when the relationship's state is `consistent_synchronized`.

An invocation example

```
stopfcmap testmapone
```

The resulting output

No feedback

Chapter 16. Host commands

The host commands enable you to work with host objects with the SAN Volume Controller.

addhostclustermember

Use the **addhostclustermember** command to add a host object to a host cluster.

Syntax

```
➡➡➡ addhostclustermember — -host —┬─ host_id_list ─┬─ hostcluster_id ─┬─┐
                                   └─ host_name_list ─┘   └─ hostcluster_name ─┘
```

Parameters

-host *host_id_list* | *host_name_list*

(Optional) Specifies the host (by ID or name) to add to the host cluster.

hostcluster_id | *hostcluster_name*

(Required) Specifies (by ID or name) that the host cluster that the host object is added to. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command adds a host object to a host cluster.

When you add a host object to a host cluster, shared mappings are created. For example, if any host mappings match a host cluster mapping that is part of the same volume on the same Small Computer System Interface (SCSI) logical unit number (LUN) - with the same I/O groups - the host cluster assumes control of the mapping (which makes it a shared mapping).

Note: A host cannot be added to a host cluster if both have their individual throttling specifications defined. However, if either the host or host cluster throttling specification is present, the command succeeds.

Any mappings that do not match the shared host cluster mappings are managed by the host as private mappings.

Note: New mappings must not conflict with a shared mapping on a host system. The command fails when there are shared mappings that conflict with the host's private mappings. This includes either:

- A volume that is being mapped - but with different SCSI LUNs
- The host that has a different volume mapped but with the same SCSI LUN as a shared mapping of the host cluster

An invocation example that adds host 0 to host cluster 4

```
addhostclustermember -host 0 4
```

The resulting output:

No feedback

An invocation example that adds hosts 0, 1, and 4 to host cluster 4

```
addhostclustermember -host 0:1:4 4
```

The resulting output:

No feedback

addhostiogrp

Use the **addhostiogrp** command to map I/O groups to an existing host object.

Syntax

```
➤➤ addhostiogrp — [ -iogrp — iogrp_list ] [ 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

```
addhostiogrp -iogrpall testhost
```

The resulting output:

No feedback

addhostport

Use the **addhostport** command to add worldwide port names (WWPNs) or Internet Small Computer System Interface (iSCSI) names from an existing host object.

Syntax

```
➤➤ addhostport — [ -saswwpn — wwpn_list ] [ -fcwwpn — wwpn_list ] [ -iscsiname — iscsi_name_list ] [ -force ] [ host_name | host_id ] ➤➤
```

Parameters

-saswwpn *wwpn_list*

(Required if you do not use **-iscsiname** or **-fcwwpn**) Specifies a list of Serial Attached SCSI (SAS) WWPNs with a 16-character hexadecimal string.

-fcwwpn *wwpn_list*

(Required if you do not use **-iscsi** or **-saswwpn**) Specifies a list of Fibre Channel (FC) WWPNs with a 16-character hexadecimal string.

-iscsi *iscsi_name_list*

(Required if you do not use **-fcwwpn** or **-saswwpn**) 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 **-fcwwpn** or **-saswwpn** 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 host bus adapter (HBA) WWPNs or iSCSI names to the specified host object. Any volumes 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 **lssasportcandidate** or **lsfcportcandidate** command.

Some HBA device drivers do not log in to the fabric until they can recognize 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 the **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 **addhostport** command. After adding the port to the host, you must create a host authentication entry using the **chhost** command.

The **addhostport** command fails if the:

- Host is mapped to a volume with more than one I/O group in the access set and the host port being added is an Internet Small Computer System Interface (iSCSI) name
- Port being added is from a host system that does not support volumes being mapped from multiple I/O groups

An invocation example

```
addhostport -saswwpn 210100E08B251DD4 host1
```

The resulting output:

No feedback

An invocation example

```
addhostport -fcwwpn 210100E08B251EE6 host1
```

The resulting output:

No feedback

An invocation example

```
addhostport -iscsi iqn.localhost.hostid.7f000001 mchost13
```

The resulting output:

No feedback

-nosite

(Optional) Resets the site value.

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 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 bus adapter (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 **lsiscsi** command does not display the authentication entry for that host. Use the **addhostiogrp** command to map the iSCSI host to at least one I/O group, and then use the **addhostport** command to add the iSCSI port into it. You must also add authentication for that host using the **chhost** command with either the **chapsecret** or **nochapsecret** parameter.

An invocation example

```
chhost -name testhostlode -mask 11111101101 hostone
```

The resulting output:

No feedback

An invocation example

```
chhost -type openvms 0
```

The resulting output:

No feedback

An invocation example

```
chhost -site site1 host3
```


The resulting output:

No feedback

chhostcluster

Use the **chhostcluster** to change the name, type, or site of a host cluster object that is part of a host cluster.

Syntax



Parameters

-name *new_name*

(Optional) Specifies a name for a host cluster object. The value must be an alphanumeric string.

-type *new_type*

(Optional) Changes the type of the all hosts in the host cluster. The values are:

- generic (the default)
- tpgs (when target port information changes, extra unit attentions are given to the host)
- openvms
- adminlun (enables virtual volumes)
- hpux
- epps
- hide_secondary

-site *new_site*

(Optional) Changes the site for all hosts in the host cluster. The default is `site0` (meaning there is no site). The value must be an alphanumeric string.

-nosite

(Optional) Resets the site value.

-force

(Optional) Specify this parameter to change the host cluster's site and also change a site for at least one host in that host cluster. (You do not have to specify this parameter if you are changing site 0.)

Remember: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information.

hostcluster_id | *hostcluster_name*

(Required) Specifies (by ID or name) the host cluster object to change. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command changes the name, type, or a site on a host cluster object.

If you assign a host to a site and it is not specified as `site0`, you must use **-force** to change it.

Important: Do not specify `-site site0`, instead use **-nosite**.

If you specify the **-site** and **-type** parameters, note that the site and type properties changed are not attributes of the host cluster object. They are properties of the individual hosts that are members of the host cluster. These properties can be modified using the **chhostcluster** command to change the attribute values on every host in the host cluster simultaneously (rather than modifying each host separately).

An invocation example that changes settings of host cluster `hostcluster0`

```
chhostcluster -name myhostcluster hostcluster0
```

The detailed resulting output:

No feedback

An invocation example that changes site of every host in host cluster 2

The hosts are currently in site 0 or in `site1`.

```
chhostcluster -site site1 2
```

The detailed resulting output:

No feedback

An invocation example that changes name of hostcluster0

All host sites are set to `site0`, which is the default. One host is not currently in site 0.

```
chhostcluster -name jvardyl -nosite -force hostcluster0
```


The detailed resulting output:

No feedback

lshost

Use the **lshost** command to generate a list with concise information about all the hosts visible to the clustered system (system) and detailed information about a single host.

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 command-line interface (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:

```
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 **lshost** command are:

- host_cluster_id
- host_cluster_name
- host_name
- host_id
- id
- iogrp_count
- name
- port_count
- site_id
- site_name
- status

Description

This command returns a concise list or a detailed view of hosts visible to the system.

For Fibre Channel (FC) ports, the `node_logged_in_count` field provides the number of nodes that the host port is logged into. For Internet Small Computer System Interface (iSCSI) ports, the `node_logged_in_count` field provides the number of iSCSI sessions from the host iSCSI qualified name (IQN).

You can map an iSCSI host to volumes that are accessible through multiple I/O groups. iSCSI hosts can access volumes that are accessible through multiple I/O groups (and single I/O groups). An iSCSI host mapped to a volume accessible through multiple I/O groups is online if it has at least one active iSCSI session with each I/O group of the access set. If volumes are not mapped to an iSCSI host, it is degraded.

The following list provides the different states for a fabric attach FC host port:

active The host port is active if all nodes with volume mappings have a login for the specified worldwide port name (WWPN) and at least one node has received SCSI commands from the WWPN within the last five minutes.

degraded

The host port is degraded if one or more nodes with volume mappings do not have a login for the specified WWPN.

inactive

The host port is inactive if all the nodes with volume mappings have a login for the specified WWPN but no nodes have seen any Small Computer System Interface (SCSI) commands from the WWPN within the last five minutes.

offline

The host port is offline if one or more input/output (I/O) groups with volume mappings do not have a login for the specified WWPN.

The following list provides the different states for a direct attach FC host port:

active The host port is active if a node has a login for the specified WWPN and the node has received SCSI commands from the WWPN within the last five minutes.

inactive

The host port is inactive if all the nodes with volume mappings have a login for the specified WWPN but no nodes have seen any Small Computer System Interface (SCSI) commands from the WWPN within the last five minutes.

offline

The host port is offline if there is no login for the specified WWPN.

If a host does not have any volume mappings it is reported as offline or inactive.

Note: The **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 **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.

The following list provides the different states for a specified `iscsiname`:

active The `iscsiname` is active if all I/O groups with volume mappings have at least one associated iSCSI session for the specified `iscsiname`.

inactive

The `iscsiname` is inactive if the host has no volume mappings but at least one iSCSI session for the specified `iscsiname` is present.

offline

The `iscsiname` is offline if one or more I/O groups with volume mappings do not have an associated iSCSI session for the specified `iscsiname`.

The following list provides the different states for `host_status`:

online The host has full connectivity. A host using just one style of connectivity is online if it uses one of these:

Fibre Attach Fibre Channel (FAFC)

Every port is active or inactive, and is logged into every online node in each I/O group in which the host has volume mappings.

Direct Attach Fibre Channel (DAFC)

The host has an active or inactive login to every node in I/O groups to which the host has volume mappings.

Internet Small Computer System Interface (iSCSI)

The host has an iSCSI session with each I/O group with which the host has volume mappings.

offline

The host has no connectivity. This might be because the host has been powered off and is not on.

Remember: If an iSCSI host is only logged into I/O groups for which it is not configured, the associated host object status is offline.

degraded

The host is not fully connected, which might be introduced by a configuration error or a hardware failure. This can cause a loss of access during any planned maintenance activity and should be corrected as soon as possible.

Remember: An iSCSI host that has no mapped volumes is degraded if it is logged in to some, but not all, of the I/O groups to which it belongs.

mask The Fiber Channel (FC) I/O ports (which exist on a node) hosts can access.

Table 81 shows the possible outputs:

Table 81. lshost output

Attribute	Description
id	Indicates the unique host ID with an alphanumeric value.
name	Indicates the unique host name with an alphanumeric string.
port_count	Indicates the number of ports.
type	Indicates the host type.
mask	Indicates the mask value with a 64-bit binary string.
iogrp_count	Indicates the number of I/O groups.
status	Indicates if the host is online or offline.
WWPN	Indicates the worldwide port name (WWPN) with a 16-character hexadecimal string.
SAS_WWPN	Indicates the serial-attached SCSI (SAS) WWPN with a 16-character hexadecimal string.
node_logged_in_count	Indicates the number of nodes the WWPN is logged into
state	Indicates the state of the SAS WWPN login: <ul style="list-style-type: none"> • offline • inactive • active
sas_wwpn_count	Indicates the number of configured SAS WWPNs.
site_id	Identifies the site ID for the host. The values are 1, 2, or blank.
site_name	Identifies the site name for the host. The value must be an alphanumeric string or blank.
host_cluster_id	Indicates the unique ID for a host cluster.
host_cluster_name	Indicates the unique name for a host cluster.

An invocation example

1shost

The resulting output:

id	name	port_count	iogrp_count	status	mapping_count	host_cluster_id	host_cluster_name	site_id	site_name
0	hostone	1	4	offline		0	var dyhost1	2	chelsea3
1	host0	1	4	degraded		1	var dyhost2	1	chelsea1
2	host1	1	4	online		2	var dyhost3	2	chelsea2

A detailed invocation example

```
1shost 0
```

The resulting output:

[illegible]

```
node_logged_in_count 1
state inactive
site_id 2
site_name chelsea2
host_cluster_id 1
host_cluster name jvardy8
```

An invocation example

```
1shost 0
```


The resulting output:

[illegible]

Ishostcluster

Use the **lshostcluster** command to generate a list with concise information about all the host clusters visible to the clustered system or detailed information about a single host 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 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 1-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.

-hostcluster *hostcluster_id | hostcluster_name*

(Required) Specifies the ID or name of the host to display information about. If you do not enter a host system ID or name, the command displays a list of all recognized host clusters and volume mappings. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command lists concise information about all the host clusters visible to the clustered system or detailed information about a single host cluster.

This table provides the attribute values that can be displayed as output view data.

Table 82. lshostcluster output

Attribute	Description
id	Indicates the host cluster ID.
name	Indicates host cluster name. The value must be an alphanumeric string of no more than 64 characters.
status	Indicates the status of the host cluster. The values are: <ul style="list-style-type: none">• online indicates that all hosts or members are online• host_degraded indicates that no hosts are offline but at least one host is degraded• host_cluster_degraded indicates that one or more hosts are offline and at least one host is online or degraded• offline indicates that all hosts are offline or there are no hosts or members in the host cluster
host_count	Indicates the number of hosts that are in the host cluster. The value must be a number from 0 - 127.
mapping_count	Indicates the number of shared mappings between the host cluster and any existing volumes. The value must be a number from 0 - 2047.
port_count	Indicates the number of host ports that are used for the host cluster mappings to any volumes. The value must be a number from 0 - 255.

A concise invocation example

```
lshostcluster
```

The detailed resulting output:

```
id name          status host_count mapping_count port_count
0  hostcluster0  online 2          1          4
```

A detailed invocation example

```
lshostcluster : hostcluster0
```

The detailed resulting output:

```
id:0
name:hostcluster0
status:online
host_count:6
mapping_count:32
port_count:12
```

lshostclustermember

Use the **lshostclustermember** command to generate a list with host information for hosts that belong to the specified host cluster.

Syntax

```
➤— lshostclustermember — [ -nohdr ] [ -delim delimiter ] [ hostcluster_id | hostcluster_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 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 1-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.

hostcluster_id | *hostcluster_name*

(Required) Specifies the ID or name of the host that is part of the host cluster. If neither a host cluster ID nor name is entered, the command displays a list of all recognized host clusters and volume mappings. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command information about all hosts that belong to the specified host cluster.

This table provides the attribute values that can be displayed as output view data.

Table 83. lshostclustermember output

Attribute	Description
host_id	Indicates the unique ID of the host cluster. The value must be a number 0 - 4095.
host_name	Indicates the host name. The value must be an alphanumeric string of no more than 64 characters.
status	Indicates the status of a host for a host cluster. The values are: <ul style="list-style-type: none">• online indicates that all hosts or members are online• host_degraded indicates that no hosts are offline but at least one host is degraded• offline indicates that all hosts are offline or there are no hosts or members in the host cluster

Table 83. *lshostclustermember* output (continued)

Attribute	Description
type	Indicates the unique ID for the site that the host cluster is in. The values are: <ul style="list-style-type: none"> • generic • hpux • tpgs • openmvs • epps • adminlun
site_id	Indicates the site ID (that the host cluster is part of). The value must be a number 0 - 3.
site_name	Indicates the site name (that the host cluster is part of). The value must be an alphanumeric string.

A concise invocation example

```
lshostclustermember
```

The detailed resulting output:

```
host_id host_name status type  site_id site_name
0      host0      online generic 1      site1
```

A detailed invocation example

```
lshostclustermember :
```

The detailed resulting output:

```
host_id:0
host_name:jlmvady
status:online
type:generic
site_id:1
site_name:jamiev12
```

lshostclustervolumemap

Use the **lshostclustervolumemap** command to display a list of volumes that are mapped to all host clusters (or to a specific host cluster).

Syntax

```
►► lshostclustervolumemap [ -nohdr ] [ -delim delimiter ] [ hostcluster_id ] [ hostcluster_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 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.

hostcluster_id | hostcluster_name

(Required) Specifies the ID or name for the host cluster being mapped to a volume. The command displays a list of all the volumes that are mapped to the specified host cluster and additionally indicates the Small Computer System Interface (SCSI) ID through which they are mapped. If you do not enter a host cluster ID or name, the command displays a list of all recognized host clusters and volume mappings. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command displays a list of volumes that are mapped to all host clusters or to a specific host cluster.

This table provides the attribute values that can be displayed as output view data.

Table 84. lshostclustervolumemap output

Attribute	Description
id	Indicates the host cluster ID. The value must be a number from 0 - 127.
name	Indicates host cluster name. The value must be an alphanumeric string of no more than 64 characters.
SCSI_id	Indicates the unique ID (volume ID) that is mapped from a host cluster to an I/O group volume. The value must be a number from 0 - 2047.
volume_id	Indicates the unique ID of the volume that is mapped to the host cluster. The value must be a number.
volume_name	Indicates the name for a volume that is mapped to a host cluster. The value must be an alphanumeric string.
volume_UID	Indicates the unique UID of a volume. The value must be an alphanumeric string.
IO_group_id	Indicates the unique ID from the I/O group that the host cluster and volume (from the mapping) are part of. The value must be a number from 0 - 3.
IO_group_name	Indicates the I/O group name. The value must be an alphanumeric string.

A concise invocation example

```
lshostclustervolumemap
```

The detailed resulting output:

```

name          SCSI_id volume_id volume_name volume_UID          IO_group_id IO_group_name
0 hostcluster0 0          0          vdisk0      60050764009900082000000000000000 0          io_grp0
0 hostcluster0 1          1          vdisk1      60050764009900082000000000000001 0          io_grp0
0 hostcluster0 2          2          vdisk2      60050764009900082000000000000002 0          io_grp0
0 hostcluster0 3          3          vdisk3      60050764009900082000000000000003 0          io_grp0
1 hostcluster1 0          4          vdisk4      60050764009900082000000000000004 0          io_grp0
1 hostcluster1 1          5          vdisk5      60050764009900082000000000000005 0          io_grp0
```

A concise invocation example

```
lshostclustervolumemap 0
```

The detailed resulting output:

id	name	SCSI_id	volume_id	volume_name	volume_UID	IO_group_id	IO_group_name
0	hostcluster0	0	0	vdisk0	60050764009900082000000000000000	0	io_grp0
0	hostcluster0	1	1	vdisk1	60050764009900082000000000000001	0	io_grp0
0	hostcluster0	2	2	vdisk2	60050764009900082000000000000002	0	io_grp0
0	hostcluster0	3	3	vdisk3	60050764009900082000000000000003	0	io_grp0

A concise invocation example

```
lshostclustervolumemap hostcluster1
```

The detailed resulting output:

id	id	name	SCSI_id	volume_id	volume_name	volume_UID	IO_group_id	IO_group_name
1	hostcluster1	0	4	vdisk4	60050764009900082000000000000004	0	io_grp0	
1	hostcluster1	1	5	vdisk5	60050764009900082000000000000005	0	io_grp0	

lshostiogrp

Use the **lshostiogrp** command to display a list the I/O groups associated with a specified host.

Syntax

```

>> 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

```
lshostiogrp -delim : hostone
```

The resulting output:

```


id:name
0:io_grp0
1:io_grp1

```

lsiscsiauth

Use the **lsiscsiauth** command to list the Challenge Handshake Authentication Protocol (CHAP) secret configured for authenticating an entity to the SAN Volume Controller clustered system (system).

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.

-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:

```
lsiscsiauth -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 **lsiscsiauth** 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 system. 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 **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 system, use the **chhost** command with the **chapsecret** parameter.

An invocation example

```
lsscsiauth
```

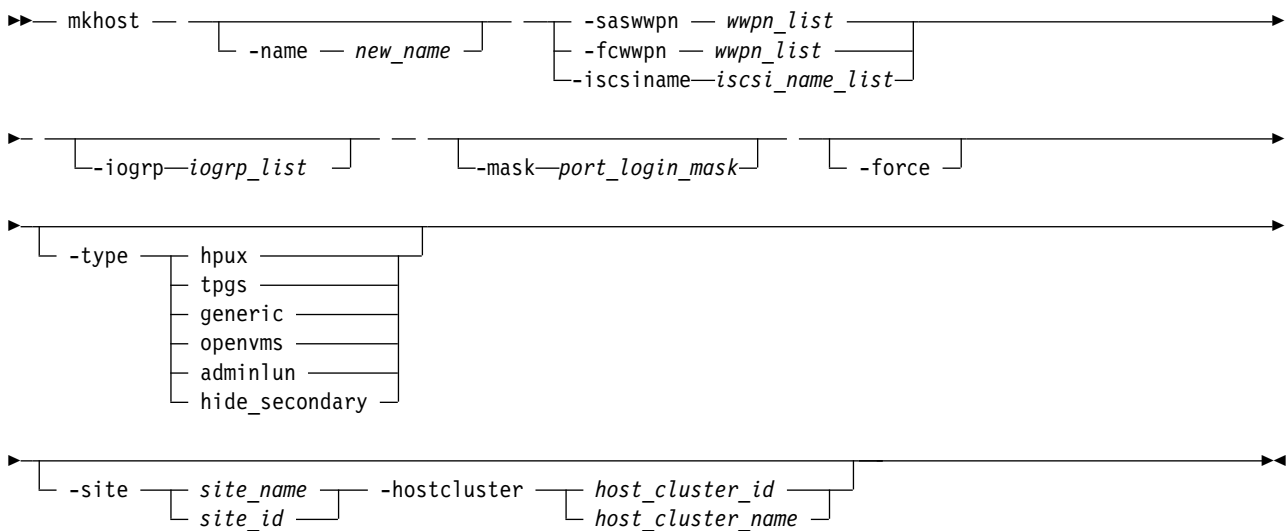
The resulting output:

type	id	name	iscsi_auth_method	iscsi_chap_secret	cluster_iscsi_auth_method	cluster_iscsi_chap_secret
host	0	mchost20	none		none	
host	1	mchost30	none		none	
host	2	mchost200	none		none	
host	3	mchost40	none		none	
host	4	mchost240	none		none	
host	5	mchost170	none		none	
host	6	mchost120	none		none	
host	7	mchost60	none		none	
host	8	mchost180	none		none	
host	9	mchost13	none		none	
host	10	newhost	none		none	

mkhost

Use the **mkhost** command to create a logical host object.

Syntax



Parameters

-name *new_name*

(Optional) Specifies a name or label for the new host object.

-saswwpn *wwpn_list*

(Required if you do not use **-iscsiname** or **-fcwwpn**) Specifies a list of Serial Attached SCSI (SAS) WWPNS with a 16-character hexadecimal string.

For additional information, see the **mkvdiskhostmap**, **lssasportcandidate**, and **lsfcportcandidate** commands.

An invocation example

```
mkhost -name hostone -saswpn 210100E08B251DD4:210100F08C262DD8 -force -mask 111111101101
```

The resulting output:

```
Host id [1] successfully created
```

An invocation example

```
mkhost -iscsiname iqn.localhost.hostid.7f000001 -name newhost
```

The resulting output:

```
Host, id [10], successfully created
```

An invocation example

```
mkhost -fcwpn 210100E08B251EE6:210100F08C262EE7 -type openvms
```

The resulting output:

```
Host, id [1], successfully created
```

An invocation example

```
mkhost -fcwpn 210100E08B251EE6 -site site1
```

The resulting output:

```
Host, id [1], successfully created
```

mkhostcluster

Use the **mkhostcluster** command to create a host cluster object.

Syntax

```
➤— mkhostcluster — -name — name —————➤
➤— -seedfromhost — host_id_list — host_name_list — -ignoreseedvolume — volume_id_list — volume_name_list —➤
```

Parameters

-name *name*

(Optional) Specifies the name of the host cluster object.

-seedfromhost *host_id_list* | *host_name_list*

(Optional) Adds the specified host to the host cluster. The host cluster mappings to existing volumes then become shared host cluster mappings. If a list of hosts is provided, the hosts are mapped to the same volume that uses the same I/O group and with the same SCSI LUN.

Note: These become shared mappings unless explicitly excluded with **-ignoreseedvolume**.

-ignoreseedvolume *volume_id_list* | *volume_name_list*

(Optional) Specifies volumes that are not part of the shared host cluster mappings. These volumes remain privately mapped to the host or hosts. You must specify **-seedfromhost** if you specify this parameter.

Description

This command is used to create a host cluster object.

Note: This command fails if any of the specified seeding hosts has an associated host throttle.

An invocation example that creates host cluster `myhostcluster` and obtains its mappings from host `myhost1`

```
mkhostcluster -name myhostcluster -seedfromhost myhost1
```

The detailed resulting output:

No feedback

An invocation example that creates host cluster `myhostcluster` that obtains its mappings from host `myhost1`

The system keeps the mapping to its boot drive (`volume_4`) private.

```
mkhostcluster -name myhostcluster -seedfromhost myhost1 -ignoreseedvolume volume_4
```

The detailed resulting output:

No feedback

An invocation example of a list being created

```
mkhostcluster -seedfromhost 1:2:3
```

The detailed resulting output:

No feedback

mkvolumehostclustermap

Use the **mkvolumehostclustermap** command to generate a new mapping between a volume and a host cluster on a clustered system. This volume is then accessible for input or output (I/O) operations to the specified host cluster.

Syntax

```
►► mkvolumehostclustermap [ -scsi scsi_num_arg ] [ -force ]
► -hostcluster [ hostcluster_id | hostcluster_name ] [ volume_id | volume_name ]
```

Parameters

-scsi scsi_num_arg

(Optional) Specifies the Small Computer System Interface (SCSI) logical unit number (LUN) ID to assign to a volume on the specified host cluster. The SCSI LUN ID is assigned to the volume on the host cluster for all I/O groups that provide access to the volume.

Note: You must use the next available SCSI LUN ID for each host in the host cluster.

-force

(Optional) This must be specified to map a volume to a host cluster and that volume is already mapped to at least one host in a different host cluster.

Remember: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information.

-hostcluster *hostcluster_id* | *hostcluster_name*

(Required) Specifies the host cluster (by ID or name) to map to the volume. The value for the ID must be a number and the value for the name must be an alphanumeric string.

volume_id | *volume_name*

(Optional) Specifies the volume by ID or name. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command generates a new mapping between a volume and a host cluster on a clustered system (system). This volume is then accessible for input or output (I/O) operations to the specified host cluster.

An invocation example that maps volume 0 to host cluster 0

```
mkvolumehostclustermap -hostcluster 0 0
```

The detailed resulting output:

No feedback

An invocation example that maps volume myvolume1 to host cluster myhostcluster and specifies SCSI LUN ID 7

```
mkvolumehostclustermap -hostcluster myhostcluster -scsi 7 myvolume1
```

The detailed resulting output:

No feedback

rmhost

Use the **rmhost** command to delete a host object.

Syntax

```
➤ 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 volumes. 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 **lsfcportcandidate** or **lssasportcandidate** command, the host objects are listed as candidate ports.

Note: This command deletes the associated host throttle if that host is removed.

Remember: This command is unsuccessful if:

- Volume protection is enabled (using the **chsystem** command)
- The host being deleted is mapped to any volume that has received I/O within the defined volume protection time period

If any mappings still exist between this host and volumes, 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

```
rmhost host_one
```

The resulting output:

No feedback

rmhostcluster

Use the **rmhostcluster** command to remove a host cluster.

Syntax

```
➤➤ rmhostcluster [ -removeallhosts | -removemappings | -keepmappings ] -hostcluster [ hostcluster_id | hostcluster_name ] ➤➤
```

Parameters

-removeallhosts

(Optional) Specifies the deletion of all hosts and the associated host cluster object.

-removemappings

(Optional) Specifies that the host cluster object being removed from the host cluster not use the host cluster's shared volume mappings. The mappings are deleted before the host cluster is deleted.

-keepmappings

(Optional) Specifies that the host cluster object removed from the host cluster retains the host cluster shared volume mappings (which become private mappings).

-hostcluster *hostcluster_id* | *hostcluster_name*

(Required) Specifies (by ID or name) that the host cluster that the host cluster object is removed from. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command removes a host cluster.

Note: This command deletes the associated host cluster throttle if that host cluster is removed.

The **-removeallhosts**, **-keepmappings**, and **-removemappings** parameters are mutually exclusive.

An invocation example that removes host cluster `hostcluster0` and any related hosts

```
rmhostcluster -removeallhosts hostcluster0
```

The detailed resulting output:

No feedback

An invocation example that removes host cluster `hostcluster0` and all mappings to volumes

```
rmhostcluster -removemappings hostcluster0
```

The detailed resulting output:

No feedback

An invocation example that removes host cluster `hostcluster0`

The hosts that are removed keep the shared mappings from the host cluster as private mappings.

```
rmhostcluster -removemappings hostcluster0
```


The detailed resulting output:

No feedback

rmhostclustermember

Use the **rmhostclustermember** command to remove a host from a host cluster object.

Syntax



Parameters

-host *host_id_list* | *host_name_list*

(Optional) Specifies (by ID or name) the hosts to remove from the host cluster.

-keepmappings

(Optional) Specifies that the host removed from the host cluster retains the host cluster's shared volume mappings. The **-keepmappings** and **-removemappings** parameters are mutually exclusive.

-removemappings

(Optional) Specifies that the host removed from the host cluster does not retain the host cluster's shared volume mappings. The **-keepmappings** and **-removemappings** parameters are mutually exclusive.

-force

(Optional) This must be specified when you remove the last host from a host cluster.

Remember: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information.

hostcluster_id | *hostcluster_name*

(Required) Specifies (by ID or name) that the host cluster that the host is removed from. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command removes a host from a host cluster object.

An invocation example that removes host 0 from host cluster 0 (and also removes the host mappings)

```
rmhostclustermember -host 0 -removemappings 0
```

The detailed resulting output:

No feedback

An invocation example that removes host myhost1 from host cluster myhostcluster while keeping the original mappings

```
rmhostclustermember -host myhost1 -keepmappings -force myhostcluster
```

The detailed resulting output:

No feedback

rmvolumehostclustermap

Use the **rmvolumehostclustermap** command to remove an existing host cluster mapping to a volume on a clustered system.

Syntax

```
➤— rmvolumehostclustermap — — -hostcluster ———— hostcluster_id —————➤
                                     | hostcluster_name |
➤— makeprivate ———— host_id_list —————➤
                       | host_name_list |
                                     | volume_id —————➤
                                     | volume_name |
```

Parameters

-hostcluster *hostcluster_id* | *hostcluster_name*

(Required) Specifies the host cluster (by ID or name) to remove from the volume mapping. The value for the ID must be a number and the value for the name must be an alphanumeric string.

-makeprivate *host_id_list* | *host_name_list*

(Optional) Specifies the host or hosts that acquire private mappings from the volume that is being removed from the host cluster. The value for the ID must be a number and the value for the name must be an alphanumeric string.

volume_id | *volume_name*

(Required) Specifies the volume by ID or name. The value for the ID must be a number and the value for the name must be an alphanumeric string.

Description

This command removes an existing host cluster mapping on a host cluster. The volume is then inaccessible for input or output (I/O) transactions from the specified host cluster.

An invocation example that moves a mapping from host cluster 0 to volume 0

```
rmvolumehostclustermap -hostcluster 0 0
```

The resulting output:

No feedback

A detailed invocation example that removes a mapping from host cluster myhostcluster and adds it to volume myvolume1

```
rmvolumehostclustermap -hostcluster myhostcluster myvolume1
```

The resulting output:

No feedback

A detailed invocation example that removes a mapping from host cluster myhostcluster and adds it to volume myvolume1

This example allows hosts myhost1 and myhost2 to acquire the private mappings from myvolume1.

```
rmvolumehostclustermap -hostcluster myhostcluster -makeprivate myhost1:myhost2 myvolume1
```

The resulting output:

No feedback

rmhostiogrp

Use the **rmhostiogrp** command to delete mappings between one or more input/output (I/O) groups and a specified host object.

Syntax

```
➤— rmhostiogrp — [ -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 host mappings.

Remember: Using the force parameter might result in a loss of access. Use it only under the direction of the your product support information.

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 **rmhostiogrp** command deletes the mappings between the list of I/O groups and the specified host object.

Remember: This command is unsuccessful if:

- Volume protection is enabled (using the **chsystem** command)
- The host I/O group being removed is mapped to any volume that has received I/O within the defined volume protection time period

If a host is defined in two I/O groups, and has access to a volume through both I/O groups, an attempt to remove the host from just one of those I/O groups fails, even with **-force** specified. To resolve this problem, do one of the following:

- Delete the host mappings that are causing the error
- Delete the volumes or the host

Note: When all I/O groups are removed from an Internet Small Computer System Interface (iSCSI) host, and you want to add an iSCSI port to the host, refer to the **addhostport** and **chhost** commands.

An invocation example

```
rmhostiogrp -iogrp 1:2 host0
```

The resulting output:

No feedback

rmhostport

Use the **rmhostport** command to delete worldwide port names (WWPNs) or Internet Small Computer System Interface (iSCSI) names from an existing host object.

Syntax

```

▶▶ rmhostport — [ -saswwpn wwpn_list ] [ -fcwwpn wwpn_list ] [ -iscsiname iscsi_name_list ] [ -force ] [ host_name | host_id ]

```

Parameters

-saswwpn *wwpn_list*

(Required if you do not use **iscsiname** or **fcwwpn**) Specifies the colon-separated list of Serial Attached SCSI (SAS) WWPNs with a 16-character hexadecimal string.

-fcwwpn *wwpn_list*

(Required if you do not use **iscsiname** or **saswwpn**) Specifies the colon-separated list of Fibre Channel (FC) WWPNs with a 16-character hexadecimal string.

-iscsiname *iscsi_name_list*

(Required if you do not use **fcwwpn** or **saswwpn**) 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 **fcwwpn** or **saswwpn** parameter.

-force

(Optional) Overrides the check that specifies all of the WWPNs or iSCSI names in the list are mapped to the host. Ports not associated with the host are ignored.

Important: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information.

host_name | *host_id*

(Required) Specifies the host name or the host ID.

Description

This command deletes the list of host-bus adapter (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.

Any volumes that are mapped to this host object are automatically unmapped from the ports.

Remember: This command is unsuccessful if:

- Volume protection is enabled (using the **chsystem** command)
- The last host port being deleted is mapped to any volume that has received I/O within the defined volume protection time period

If multiple hosts are mapped to the same active volume, the host port removal is allowed if the host is offline. This allows for the removal of ports from hosts that might be part of the same system.

List the candidate Fibre Channel (FC) or serial-attached SCSI (SAS) ports by issuing the **lsfcportcandidate** or **lsasportcandidate** command. A list of the ports that are available to be added to host objects is displayed. To list the WWPNs that are currently assigned to the host, issue the following:

```
lshost hostobjectname
```

where *hostobjectname* is the name of the host object.

Add the new ports to the existing host object by issuing the following command:

```
addhostport -fcwwpn one or more existing WWPNs  
separated by : hostobjectname/ID
```

where one or more existing WWPNs separated by colon (:) and *hostobjectname/id* is the name or ID of the host object.

Remove the old ports from the host object by issuing the following command:

```
rmhostport -fcwwpn one or more existing WWPNs  
separated by : hostobjectname/ID
```

where one or more existing WWPNs separated by colon (:) corresponds with those WWPNs that are listed in the previous step. Any mappings that exist between the host object and volumes are automatically applied to the new WWPNs.

An invocation example

```
rmhostport -saswwpn 210100E08B251DD4 host1
```

The resulting output:

No feedback

An invocation example

```
rmhostport -fcwwpn 210100E08B251EE6 host1
```

The resulting output:

No feedback

An invocation example

```
rmhostport -iscsiname iqn.localhost.hostid.7f000001 mchost13
```

The resulting output:

No feedback

Chapter 17. Information commands

The information commands enable you display specific types of SAN Volume Controller information.

These commands return no output but exit successfully when there is no information to display.

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.

ls2145dumps (Deprecated)

The **ls2145dumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

lsconfigdumps (Discontinued)

The **lsconfigdumps** command is discontinued. Use the **lsdumps** instead.

lssshkeys (Discontinued)

Attention: The **lssshkeys** command is discontinued. Use the user management commands to configure remote authentication service and manage users and user groups on the cluster.


Chapter 18. Key server commands

Key server commands are used to create, change, or list details about key servers.

chkeyserver

Use the **chkeyserver** command to change the attributes for a key server object.

Syntax



Parameters

-ip *ip_address*

(Optional) Specifies the key server's IP address. The value must be in the form of a standard Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) address.

-port *port*

(Optional) Specifies the key server's TCP/IP port. The value must be a number 1 - 65535. The default value is the same as the default listening port used for key servers of the currently enabled type.

-sslcert *certificate_file*

(Optional) Specifies the key server's self-signed certificate. The value must be a file path string.

-nossllcert

(Optional) Specifies the removal of the key server's self-signed certificate.

-name

(Optional) Specifies the key server object name. The value must be an alphanumeric string.

-primary

(Optional) Specifies the primary key server.

object_id | *object_name*

(Required) Specifies the object name or ID that you want to modify.

Description

This command changes the attributes for a key server object.

An invocation example

```
chkeyserver -primary varyd2
```

The resulting output:

No feedback

An invocation example

```
chkeyserver -nosslcert
```


The resulting output:

No feedback

chkeyserverisklm

Use the **chkeyserverisklm** command to change the system-wide IBM Security Key Lifecycle Manager key server configuration.

Syntax



Parameters

-devicegroup *device_group*

(Optional) Specifies a specific device group that the system uses with a key server. The value must be an alphanumeric string no more than 16 characters long.

Note: The specified device name must begin with a letter (not a number) and cannot contain an underscore.

-sslcert *certificate_file*

(Optional) Specifies the certificate authority (CA) certificate for the key server. The value must be an alphanumeric string.

-nosslcert

(Optional) Specifies that the CA certificate on the key server be deleted.

-enable

(Optional) Enables the specified key server type.

-disable

(Optional) Disables the specified key server type.

Important: Do not specify **-disable** with other parameters.

Description

This command changes the system-wide IBM Security Key Lifecycle Manager key server configuration.

An invocation example

```
chkeyserverisklm -devicegroup JVAR_IBRA -sslcert /dumps/CA_certificate.pem -enable
```

The resulting output:

No feedback

An invocation example

```
chkeyserverisklm -nosslicert
```

The resulting output:

No feedback

Iskeyserver

Use the **Iskeyserver** command to display the key servers that are available to the clustered system (system).

Syntax

```
➤— Iskeyserver — [ -nohdr ] [ -delim delimiter ] [ object_id | object_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 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 1-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 object name or ID that you want to display details for.

Description

This command displays all key servers that are available to the system.

This table provides the attribute values that can be displayed as output view data.

Table 85. Iskeyserver output

Attribute	Description
id	Indicates the key server ID. The value must be a number.
name	Indicates the key server name. The value must be an alphanumeric string.
status	Indicates the key server status type. The values are: <ul style="list-style-type: none">• online• degraded• offline
err_seq_num	Indicates the event log sequence number of the highest priority problem that affects the key server. The value must be a number.

Table 85. *lskeyserver* output (continued)

Attribute	Description
IP_address	Indicates the key server Internet Protocol (IP) address. The value must be in the format of a standard Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) address.
port	Indicates the key server TCP/IP port.
type	Indicates the key server type. The value must be <code>isklm</code> .
primary	Indicates if the server is a primary server. The values are <code>yes</code> or <code>no</code> .
cert_set	Indicates whether a certificate exists for this key server object. The values are <code>yes</code> or <code>no</code> .
certificate	Indicates a human-readable description of the SSL certificate. The value reads 0 fields if there is no certificate.

An invocation example

```
lskeyserver
```

The resulting output:

```
id name      status  IP_address      port  type    primary  cert_set
0  isklm_primary online  10.0.1.54      8709  isklm    yes      yes
1  isklm_backup online  10.0.1.55      8709  isklm    no       yes
2  keyserver2  offline 0:0:0:0:ffff:a00:138 1234  isklm    no       no
3  keyserver3  offline 0:0:0:0:ffff:a00:139 1234  isklm    no       no
```

An invocation example

```
lskeyserver 0
```

The resulting output:

```
id 0
name keyserver0
status online
err_seq_num
IP_address 10.0.1.54
port 8709
type isklm
primary yes
certificate 0 fields
```

lskeyserverisklm

Use the **lskeyserverisklm** command to display the system-wide IBM Security Key Lifecycle Manager key server configuration.

Syntax

```
➤➤ lskeyserverisklm — [ -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 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 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 1-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 system-wide IBM Security Key Lifecycle Manager key server configuration.

This table provides the attribute values that can be displayed as output view data.

Table 86. Iskeyserverisklm output

Attribute	Description
status	Indicates the key server status type. The values are: <ul style="list-style-type: none">• disabled• enabled_inactive• prepared• enabled_active
device_group	Indicates the device group. The value is a 16-character alphanumeric string.
certificate	Indicates a human readable description of the server SSL certificate that is generated by the system.

An invocation example

```
lskeyserverisklm
```

The detailed resulting output:

```
status enabled_active
device_group VARDY_SYSTEM
certificate 58 fields
Data:
  Version: 3 (0x2)
  Serial Number: 1431938814 (0x5559a6fe)
  Signature Algorithm: sha256WithRSAEncryption
  Issuer: C=GB, L=Hursley, O=IBM, OU=SSG, CN=2145/emailAddress=support@ibm.com
  Validity
    Not Before: May 18 08:46:54 2015 GMT
    Not After : May 14 08:46:54 2030 GMT
  Subject: C=GB, L=Hursley, O=IBM, OU=SSG, CN=2145/emailAddress=support@ibm.com
  Subject Public Key Info:
    Public Key Algorithm: rsaEncryption
    Public-Key: (2048 bit)
    Modulus:
      00:de:1c:70:c2:91:87:3c:6a:92:91:f7:d9:a3:5b:
      05:e6:91:f1:87:c1:25:38:61:ad:4d:d9:26:19:7b:
      9e:61:a5:fd:b1:d1:eb:d1:e4:a8:78:21:75:58:80:
      4a:5c:dd:5e:6c:8b:1b:de:57:f9:d5:1f:71:92:3e:
      78:d5:a4:75:1e:11:b2:62:18:52:0f:4d:32:a8:fd:
      2b:16:4f:42:d1:d6:70:af:86:eb:fe:a1:ab:bc:66:
      8a:44:bc:e0:36:53:77:96:2f:74:7d:95:33:79:c2:
      59:5e:e1:43:50:da:43:25:c4:5d:3a:ac:d7:82:ad:
      34:d5:ba:4c:52:4a:c0:81:3a:ad:e8:33:fe:4f:be:
      e8:47:fa:5b:1f:dd:d8:9e:3b:44:a6:b6:b9:43:d2:
      d4:45:8e:cb:5b:bb:10:5b:c9:30:68:2c:30:b6:e4:
      ea:59:6d:a2:37:a7:13:77:28:1d:13:68:58:7b:dd:
```

```

90:d6:a8:81:7b:79:9f:1e:e4:a7:67:1b:7b:c5:b4:
90:dc:6b:d4:1f:7e:e9:e3:7b:ac:26:59:11:f1:99:
34:f0:6a:50:41:76:ad:a3:30:74:8f:8f:f5:ed:1e:
21:77:ff:51:90:1b:83:fb:04:f0:62:3d:71:17:a5:
ab:44:e8:bc:b0:82:0d:af:af:ae:68:5a:cf:e3:c8:
a9:53
Exponent: 65537 (0x10001)
X509v3 extensions:
  X509v3 Basic Constraints:
    CA:FALSE
  Netscape Comment:
    OpenSSL Generated Certificate
  X509v3 Subject Key Identifier:
    87:66:33:16:61:7A:8E:CA:B4:BA:78:7B:56:56:8A:9D:C5:96:80:76
  X509v3 Authority Key Identifier:
    keyid:87:66:33:16:61:7A:8E:CA:B4:BA:78:7B:56:56:8A:9D:C5:96:80:76

```

```

Signature Algorithm: sha256WithRSAEncryption
56:b1:5d:59:11:ae:7b:6e:29:cc:1f:a8:75:77:d2:65:d6:88:
75:8e:b9:cd:d6:71:ac:7e:89:8c:65:68:36:a8:28:97:88:36:
42:da:a4:58:9b:c6:ce:c1:56:c9:0e:c5:ce:e7:01:74:d0:66:
d0:4d:d3:0f:84:53:f6:e5:89:8e:44:6d:70:13:45:9c:21:91:
50:f4:b0:b7:cc:cb:18:e8:d7:b3:38:b4:f5:5d:36:51:8c:7e:
52:d4:24:0f:1f:2e:0a:b4:b6:9b:cb:23:43:6c:16:a2:a5:de:
84:8a:0d:28:3c:d9:3d:5d:a4:52:44:28:90:98:a6:26:a9:c9:
87:6c:27:3f:ef:09:5f:9d:0b:40:8d:07:64:ee:33:d9:40:47:
98:02:10:58:2b:54:33:d9:37:69:d4:13:e6:0d:ec:46:26:b1:
c1:c5:15:7c:8d:89:26:f7:95:d9:2f:d9:33:8c:f0:1a:dc:08:
19:eb:18:16:51:30:a3:c0:ee:be:86:7d:3d:91:61:d5:99:bf:
5e:19:b9:89:72:e1:4c:ea:5e:2b:90:ce:ce:75:83:e0:c9:14:
83:21:21:e0:f8:28:94:90:71:e6:13:ca:97:8c:e3:58:b9:0c:
62:03:e5:1c:1b:6c:dd:c3:60:48:d4:78:24:8e:22:34:78:32:
fe:45:ee:36

```

```

-----BEGIN CERTIFICATE-----
MIIDzTCCArWgAwIBAgIEVmm/jANBgkqhkiG9w0BAQsFADBgMQswCQYDVQQGEwJH
QjEQMA4GA1UEBwwHSHVyc2xleTEEMMAoGA1UECgwDSUJNMQwwCgYDVQQLDANTU0cx
DTALBgNVBAMBDIxNDUxHjAcBgkqhkiG9w0BCQEWd3N1cHBvcnRAaWJtLmNvbTAe
Fw0xNTA1MTgwODQ2NTRaFw0zMDA1MTQwODQ2NTRaMGoxCzAJBgNVBAYTAkdCMRAw
DgYDVQHDAdIdXZjbGV5MQwwCgYDVQQKDANJQk0xDDAKBgNVBAsMAINTRzENMASG
A1UEAwEMjE0NTEeMBwGCsqSgSIb3DQEJARYPc3VwcG9ydEBpYm0uY29tMIIjAN
BgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEA3hxwppGHPGqSkffZo1sF5pHxh8E1
0GGtTdkmGXueYaX9sdHr0eSoeCF1WIBKXN1ebIsb31f51R9xkj541aR1HhGyYhS
D00yqP0rFk9C0dZwr4br/qGrvGaKRLzgN1N31i90fZUzecJZXuFDUNpDjCrdQqzX
gg001bpMUKrAgTqt6DP+T77oR/pbH93YnjtEpra5Q9LURY7LW7sQW8kwaCwwtuTq
WW2iN6cTdygdE2hYe92Q1qiBe3mfHuSnZxt7xbSQ3GvUH37p43usJlkr8Zk08GpQ
QXatozB0j4/17R4hd/9RkBuD+wTwYj1xF6WrR0i8sIINr6+uaFrP48ipUwIDAQAB
o3sweTAJBGNVHRMEAjAAMCwGCWCGSAGG+EIBDQqFh1PcGVuU1NMIEdlbnVpYXR1
ZCBZDZXJ0aWZpY2F0ZTAdBgNVHQ4EFgQUh2YzFmF6jsq0unh7V1aKncWgHYwHwYD
VR0jBBgwFoAUh2YzFmF6jsq0unh7V1aKncWgHYwDQYJKoZIhvcNAQELBQADggEB
AFaxXVkrRntuKcwfqHV30mXWiHW0uc3Wcax+iYx1aDaoKJeInkLapFibxs7BVsk0
xc7nAXTQZtBN0w+EU/b1iY5EbXATRZwhkVD0sLfMyxjo17M4tPvDN1GMf1LUJA8f
Lgq0tpvLI0NsFqK13oSKDSg821dpfJEKJCYpiapyYdsJz/vCV+dC0CNB2TuM91A
R5gCEfgrVDPZN2nUE+YN7EYmscHFFXyNiSb31dkv2TOM8BrcCBnrGBZRMKPA7r6G
fT2RYdWZv14ZuY1y4UzqXiuQzs51g+DJFIMhIeD4KJSQceYTypeM41i5DGID5Rwb
bN3DYEjUeCS0IjR4Mv5F7jY=
-----END CERTIFICATE-----

```

mkkeyserver

Use the **mkkeyserver** command to create a key server object.

Syntax

```
➤ mkkeyserver — — -ip ip_address — [ -port port ] [ -sslcert certificate_file ]
➤ [ -name ] [ -primary ]
```

Parameters

- ip** *ip_address*
(Required) Specifies the key server's IP address. The value must be in the form of a standard Internet Protocol version 4 (IPv4) or Internet Protocol version 6 (IPv6) address.
- port** *port*
(Optional) Specifies the key server's TCP/IP port. The value must be a number 1 - 65535. The default value is the same as the default listening port used for key servers of the currently enabled type.
- sslcert** *certificate_file*
(Optional) Specifies the key server's self-signed certificate. The value must be a file path string.
- name**
(Optional) Specifies the key server object name. The value must be an alphanumeric string.
- primary**
(Optional) Specifies the primary key server.

Description

This command creates a key server object.

The first key server object must be the primary object, which is created by specifying the **-primary** parameter. If key management is enabled, you must use the primary key server object to create keys.

When a key server object is created, it is automatically validated. If the validation is not successful, the command fails and an error message is displayed.

An invocation example

```
mkkeyserver -ip 10.0.1.54 -sslcert /tmp/isklm_public_server_cert.pem -primary
```

The resulting output:

```
Key Server, id [0], successfully created
```

rmkeyserver

Use the **rmkeyserver** command to remove a key server object.

Syntax

```
➤ rmkeyserver — [ object_id ] [ object_name ]
```

Parameters

- object_id* | *object_name*
(Required) Specifies the object name or ID to remove.

Description

This command removes a key server object.

An invocation example

```
rmkeyserver 1
```

The resulting output:

No feedback

testkeyserver

Use the **testkeyserver** command to test key server objects.

Syntax

```
►► testkeyserver — [ object_id | object_name ] ◄◄
```

Parameters

object_id | *object_name*

(Required) Specifies the object name or ID to be validated.

Description

This command tests key server objects.

An invocation example

```
testkeyserver 0
```

The resulting output:

The key server task completed successfully.

Chapter 19. Livedump commands

Use the livedump commands to manage the node livedumps on SAN Volume Controller.

cancellivedump

Use the **cancellivedump** command to cancel a live dump.

Syntax

```
➤— cancellivedump — node_name —————➤  
                    |  
                    node_id
```

Parameters

node_name|node_id
(Required) Identifies the node name or ID.

Description

Use this command if you issue a **preplivedump** command, but then decide not to issue a **triggerlivedump** command. This releases the resources you allocated for the livedump. This event is recorded in the node trace (.trc) file. For this command to succeed, the node must be in a livedump prepared state.

An invocation example

```
cancellivedump node1
```

The resulting output:

No feedback

lslivedump

Use the **lslivedump** command to query the livedump state of a node.

Syntax

```
➤— lslivedump — node_name —————➤  
                |  
                node_id
```

Parameters

node_name|node_id
(Required) Identifies the node name or ID.

Description

You can issue this command repeatedly to determine if a live dump is in progress for the node. Table 87 on page 494 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 87. *lslivedump* outputs

Attribute	Description
inactive	The node has no livedump activity.
prepared	The node is ready to be triggered.
dumping	The node is writing the dump file.

An invocation example

```
lslivedump node1
```

The resulting output:

```
status
prepared
```

prelivedump

Use the **prelivedump** command to reserve the system resources that are required for livedump.

Syntax

```
➤ — prelivedump — [ node_name | node_id ] —➤
```

Parameters

node_name | *node_id*
(Required) Identifies the node name or ID.

Description

You can prepare more than one node for livedump at a time by issuing the **prelivedump** command consecutively. However, you can only trigger one livedump at a time, with an automatic lag time of 30 seconds between each trigger event. This helps maintain node stability.

You can issue multiple **prelivedump** commands on the same node; however, only a **prelivedump** command followed by a **triggerlivedump** command results in output.

Because the livedump resource allocation can take time to execute, you can issue this command to prepare the livedump but trigger it at a later time. This command times out after 60 seconds. The prelivedump event is located in the node trace (.trc) file.

An invocation example

```
prelivedump node1
```

The resulting output:

```
No feedback
```

triggerlivedump

Use the **triggerlivedump** command to capture the metadata that you want to dump, and write the dump file to the internal disk on the node.

Syntax

►► triggerlivedump — 

Parameters

node_name|node_id

Identifies the node name or ID.

Description

You can issue this command to trigger a **livedump** command. Only one **triggerlivedump** action can be in progress at one time, with an automatic lag time of 30 seconds between each trigger event. The node must have a livedump state of prepared for this command to succeed. Output is recorded in the node trace (.trc) file.

After you issue the **triggerlivedump** command, the command captures data and returns you to the CLI interface so that you can issue additional commands. While you issue additional commands, the livedump disk file is written to the disk in the background, and the livedump state shows as dumping. After the write is complete, the state shows as inactive.

An invocation example

```
triggerlivedump node1
```

The resulting output:

No feedback

Chapter 20. Managed disk commands


The managed disk commands enable you to work with managed disk options with the SAN Volume Controller.

If the clustered system (system) 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 system only deletes the MDisk from the list if the MDisk is offline and it has a mode of unmanaged (it does not belong to a storage pool).

addmdisk

Use the **addmdisk** command to add one or more managed disks to an existing storage pool.

Syntax



Parameters

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Required) Specifies one or more managed disk IDs or names to add to the storage pool.

-tier *tier0_flash* | *tier1_flash* | *tier_enterprise* | *tier_nearline*

(Optional) Specifies the tier of the MDisk or MDisk being added. Unless otherwise specified, the current tier value associated with the MDisk is retained. The values are:

tier0_flash

Specifies a *tier0_flash* hard disk drive or an external MDisk for the newly discovered or external volume.

tier1_flash

Specifies an *tier1_flash* (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

tier_enterprise

Specifies a *tier_enterprise* hard disk drive or an external MDisk for the newly discovered or external volume.

tier_nearline

Specifies a *tier_nearline* hard disk drive or an external MDisk for the newly discovered or external volume.

The default value for a newly discovered unmanaged MDisk is *enterprise*. You can change this value by using the **chmdisk** command.

The tier of external managed disks is not detected automatically and is set to enterprise. If the external managed disk is made up of flash drives or nearline Serial Attached SCSI (SAS) drives and you want to use Easy Tier, you must either specify the tier when adding the managed disk to the storage pool or use the **chmdisk** command.

mdisk_group_id | mdisk_group_name

(Required) Specifies the ID or name of the storage pool to add the disks to. When an MDisk is added, the warning threshold for the storage pool is automatically scaled.

Description

This command adds the managed disks that you specify to the storage pool.

If there are no MDisks in the storage pool, the site of the MDisk being added must be well-defined. If there are MDisks in the storage pool, the site information for an MDisk being added to a storage pool with HyperSwap or stretched topology system must match the topology of other MDisks in the storage pool.

Remember: This command cannot be used for child pools.

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 storage pool cannot be added to another storage pool until they have been deleted from their current storage pool. You can delete a managed disk from a storage pool under the following circumstances:

- If the managed disk does not contain any extents in use by a volume
- If you can first migrate the extents in use onto other free extents within the storage pool.

Remember: Do not include an Mdisk in an storage pool if it can only be used in image mode.

If the system has I/O groups that are not capable of encryption, you cannot add the MDisk if the MDisk group has an encryption key and the MDisk is not self-encrypting.

An invocation example

```
addmdisk -mdisk mdisk13:mdisk14 -tier tier_nearline Group0
```

The resulting output:

No feedback

applymdisksoftware (Discontinued)


Attention: The **applymdisksoftware** command has been discontinued. Use the **applydrivesoftware** command to update drives.

chmdisk

Use the **chmdisk** command to modify the name or IBM Easy Tier settings for a managed disk (MDisk).

Syntax

```
➤ chmdisk — [ -name — new_name_arg ] [ -tier — tier0_flash | tier1_flash | tier_enterprise | tier_nearline ]
```



Parameters

-name *new_name_arg*

(Optional) Specifies the new name to be applied to the managed disk.

-tier *tier0_flash* | *tier1_flash* | *tier_enterprise* | *tier_nearline*

(Optional) Specifies the new tier of the MDisk. The values are:

tier0_flash

Specifies a *tier0_flash* hard disk drive or an external MDisk for the newly discovered or external volume.

tier1_flash

Specifies an *tier1_flash* (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

tier_enterprise

Specifies a *tier_enterprise* hard disk drive or an external MDisk for the newly discovered or external volume.

tier_nearline

Specifies a *tier_nearline* hard disk drive or an external MDisk for the newly discovered or external volume.

-easytierload *default* | *low* | *medium* | *high* | *very_high*

(Optional) Specifies the Easy Tier load (amount) to place on a non-array MDisk within its tier.

If Easy Tier is either overusing or under-utilizing a particular MDisk, modify the *easy_tier_load* value to change the load size.

Note: Specifying *default* returns the performance capability to the value used by the system. Specify *very_high* only if the MDisk tier is *ssd*.

-encrypt *yes* | *no*

(Optional) Specifies whether the MDisk is encrypted using its own encryption resources. The values are *yes* or *no*.

Important: If you use SAN Volume Controller in front of an encrypted Storwize V7000 system, you should upgrade Storwize V7000 before applying encryption to your Storwize V7000 system.

If you apply encryption to your system, you must identify the encrypted MDisk before applying the encryption. If you specify *chmdisk -encrypt*, the setting is permanent in SAN Volume Controller no matter what Storwize V7000 says.

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.

Do not use the **-encrypt** parameter if one of the MDisk group's has an encryption key, parent pool, and child pools. Use **chmdisk** for existing self-encrypting MDisk before starting any migration. If an MDisk is self-encrypting, the encrypted property defaults to what is reported.

An invocation example

```
chmdisk -tier tier0_flash mdisk13
```

The resulting output:

No feedback

An invocation example

```
chmdisk -tier tier_nearline mdisk0
```

The resulting output:

MDisk Group, id [13], successfully created

An invocation example

```
chmdisk -easytierload high mdisk0
```

The resulting output:

MDisk Group, id [13], successfully created

An invocation example

```
chmdisk -name my_first_mdisk -encrypt yes 0
```

The resulting output:

MDisk Group, id [0], successfully changed

detectmdisk

Use the **detectmdisk** command to manually rescan the Internet Small Computer Systems Interface (iSCSI) or Fibre Channel (FC) network for any new managed disks (MDisks) that might have been added, and to rebalance MDisk access across all available controller device ports.

Syntax

►► — detectmdisk — -scope — *scope_id* —————►◄

Parameters

-scope *scope_id*

(Optional) Specifies the domain index. The value must be a number from 0 to 6. For example, the value 0 indicates FC and 6 indicates iSCSI.

Description

This command causes the clustered system (system) to rescan the FC or iSCSI network. The rescan discovers any new MDisks that have been added to the system 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 system automatically detects disks when they appear on the network. However, some FC 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 system has not detected it, you might need to run this command before the system detects the new disks.

When back-end controllers are added to the FC SAN and are included in the same switch zone as a system, the system 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 system might be unaware of these configuration changes. Run this command to rescan the FC or iSCSI network and update the list of unmanaged MDisks.

Note: The automatic discovery that is performed by the system does not write to an unmanaged MDisk. Only when you add an MDisk to a storage pool, or use an MDisk to create an image mode volume, is the storage actually used.

To identify the available MDisks, issue the **detectmdisk** command to scan the FC or iSCSI network for any MDisks. When the detection is complete, issue the **lsmdiskcandidate** command to show the unmanaged MDisks; these MDisks have not been assigned to a storage pool. Alternatively, you can issue the **lsmdisk** 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 **detectmdisk** command to force the SAN Volume Controller to update its configuration and accept the changes to the controller ports.

Note: Only issue the **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

```
detectmdisk
```

The resulting output:

```
No feedback
```

An invocation example

```
detectmdisk -scope 1
```

The resulting output:

```
No feedback
```

dumpallmdiskbadblocks

Use the **dumpallmdiskbadblocks** command to dump bad block counts to a dump file used by the fix procedures and the satask **snap** command.

Syntax

►► — `dumpallmdiskbadblocks` — ————— ►►

Parameters

None

Description

Use the **dumpallmdiskbadblocks** command to dump bad block counts to a readable ASCII dump file for use by fix procedures and the satask **snap** command. The output contains bad blocks for which an error log has been raised.

Use **lsdumps -prefix /dumps/mdisk** to list the output files. Use **cleardumps -prefix /dumps/mdisk** to clear the output files.

The maximum number of dump files is 20.

An invocation example

```
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

Use the **dumpmdiskbadblocks** command to write the bad block counts and locations that are on a specified MDisk to a dump file for use by fix procedures.

Syntax

►► — `dumpmdiskbadblocks` — ——— `object_id` ————— ►►
 └ `object_name` ─┘

Parameters

object_id | *object_name*

(Required) Specifies the MDisk for which you need to dump the bad block record table.

Description

Use the **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 fix procedures. The output consists of bad blocks for which an event log has been raised.

Use **lsdumps -prefix /dumps/mdisk** to list the output files. Use **cleardumps -prefix /dumps/mdisk** to clear the output files.

The reported event log sequence numbers correspond to the first event seen in the bad block record, which is a 512-block region.

- If there are multiple event logs in the same region, the earliest event sequence is used.
- If there are event logs of different types in the same region, event sequence numbers for bad blocks caused by medium errors on RAID member drives 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

```
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

Event log sequence number: 1

Start LBA: 0x5678567812341234

Length: 4

Event 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 clustered system (system).

Syntax

```
➤ 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 system.

Description

The specified managed disk is included in the system.

You might exclude a disk from the system 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 system.

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 storage pool, issuing an include command for this MDisk results in the MDisk record being deleted from the system.

An invocation example

```
includemdisk mdisk5
```

The resulting output:

No feedback

lsmdisk

Use the **lsmdisk** command to display a concise list or a detailed view of managed disks (MDisks) visible to the clustered system (system). It can also list detailed information about a single MDisk.

Syntax

```
➤ lsmdisk — [ -filtervalue — attribute=value — [ -unit [ b | kb | mb | gb | tb | pb ] ] ] ➤
```



```
➤ [ -filtervalue? ] [ -nohdr ] [ -bytes ] [ -delim — delimiter ] ➤
```



```
➤ [ object_id | object_name ] ➤
```

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. Use the **unit** parameter to interpret the value for size or capacity.

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:

```
lsmdisk -filtervalue "name=md*"
```

-unit **b** | **kb** | **mb** | **gb** | **tb** | **pb**

(Optional) Specifies the data units for the **-filtervalue** parameter.

Note: **-unit** must be used with **-filtervalue**.

-filtervalue?

(Optional) Includes all of the valid filter attributes in the report. The following filter attributes are valid for the **lsmdisk** command:

- block_size
- capacity
- controller_id
- controller_name
- ctrl_LUN_#
- easy_tier_load
- id
- max_path_count
- mode
- mdisk_grp_id
- mdisk_grp_name
- name
- path_count
- quorum_index
- site_id
- site_name
- status
- tier
- UID

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 system. Table 88 provides the potential output for MDiskS.

Table 88. 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 flash drive 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 MDiskS) or by the user:</p> <ul style="list-style-type: none">• tier0_flash• tier1_flash• tier_enterprise• tier_nearline <p>Note: You can change this value using the chmdisk command.</p>
easy_tier_load	<p>This value controls Easy Tier settings, and is either blank (for arrays) or one of the following values (for MDiskS):</p> <ul style="list-style-type: none">• low• medium• high• very_high

Table 88. MDisk output (continued)

Attribute	Values
raid_status	<p>offline Array is offline on all nodes</p> <p>degraded Array has deconfigured or offline members; the array is not fully redundant</p> <p>syncing Array members are all online, the array is syncing parity or mirrors to achieve redundancy</p> <p>initting Array members are all online, the array is initializing; the array is fully redundant</p> <p>online Array members are all online, and the array is fully redundant</p>
raid_level	The RAID level of the array (RAID0, RAID1, RAID5, RAID6, RAID10).
redundancy	The number of how many member disks that 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> <p>exact All populated members have exact capability match, exact location match</p> <p>yes All populated members have at least exact capability match, exact chain, or different enclosure or slot</p> <p>no Anything else not included for yes or exact</p>
site_id	Indicates the site value for the MDisk. This numeric value is 1, 2, 3 or blank.
site_name	Indicates the site name for the MDisk. This is an alphanumeric value or is blank.
fabric_type	<p>Indicates a Fibre Channel (FC), SAS, or another type of MDisk.</p> <ul style="list-style-type: none"> • fc indicates an MDisk from an FC controller • sas_direct indicates an MDisk from an SAS direct-attached controller • iscsi indicates an iSCSI controller
encrypt	<p>Indicates whether the data stored on the MDisk group is encrypted or not encrypted. The values are:</p> <ul style="list-style-type: none"> • yes if the pool has an encryption key • yes if the pool does not have an encryption key but it contains an MDisk or MDisk that are encrypted • no if the pool does not have an encryption key but the pool contains at least one unencrypted MDisk • Blank if the pool doesn't have an encryption key and the pool has no MDisk • yes • no
distributed	Indicates whether the array is distributed. The values are yes or no.
drive_class_id	Indicates the drive class that makes up this array. If -allowsuperior was used during array creation, the lowest used drive class ID is displayed. This value is blank for traditional arrays.
drive_count	Indicates the total width of the array, including rebuild areas. The value is a number from 4 to 128. The minimum value for RAID-6 and RAID-10 arrays is 6.

Table 88. MDisk output (continued)

Attribute	Values
stripe_width	Indicates the width of a single unit of redundancy within a distributed set of drives. The values are: <ul style="list-style-type: none"> Any number from 3 - 16 for RAID-5 arrays Any number from 4 - 16 for RAID-6 arrays An even number from 2 - 16 for RAID-10 arrays
rebuild_areas_total	Indicates the total number of rebuild areas set at array creation time. These rebuild areas provide performance but no capacity. The value is 1 - 4 for distributed array RAID-5 and RAID-6, and the value is 2 - 4 for distributed array RAID-10 (the value is blank for traditional arrays).
rebuild_areas_available	Indicates the number of remaining build areas within the set of arrays. The value is 1 - 4 for distributed array RAID-5 and RAID-6, and the value is 2 - 4 for distributed array RAID-10 (the value is blank for traditional arrays).
rebuild_areas_goal	Indicates the rebuild areas threshold (minimum limit) at which point the array logs an error. The value is 1 - 4 for distributed array RAID-5 and RAID-6, and the value is 2 - 4 for distributed array RAID-10 (the value is blank for traditional arrays).
dedupe	Indicates that dedupe is enabled. If dedupe is enabled, duplicate copies of repeating data are compressed or removed.
ctrl_WWNN	Indicates the control worldwide node name (WWNN).
preferred_WWPN	Indicates the preferred worldwide port name (WWPN).
active_WWPN	Indicates the active WWPN.
preferred_iscsi_port_id	Indicates the preferred I/O port identifier, which has the same value as the preferred_WWPN value in the Fibre Channel (FC) domain. The Internet Small Computer System Interface (iSCSI) port ID value is displayed, but the value is blank for non-iSCSI domains. This value must be a numeric value that can range from 0 - 1023.
active_iscsi_port_id	Indicates the active I/O port identifier, which has the same value as the active_WWPN value in the FC domain. The Internet Small Computer System Interface (iSCSI) port ID value is displayed, but the value is blank for non-iSCSI domains. This value must be a numeric value that can range from 0 - 1023.

Note: The automatic discovery performed by the system does not write anything to an unmanaged MDisk. It is only when you add an MDisk to a storage pool, or use an MDisk to create an image mode volume, that the system uses the storage.

To see which MDisks are available, issue the **detectmdisk** command to manually rescan the Fibre Channel or iSCSI network for any new MDisks. Issue the **lsmdiskcandidate** command to show the unmanaged MDisks. These MDisks have not been assigned to a storage pool.

Notes:

1. A SAN Volume Controller connection from a node or node canister port to a storage controller port for a single MDisk is a path. The Mdisk *path_count* value is the number of paths currently being used to submit input/output (I/O) to this MDisk.
2. The MDisk *max_path_count* value is the highest value *path_count* has reached since the MDisk was last fully online.
3. The *preferred_WWPN* is one of the World Wide Port Names (WWPNs) the storage controller has specified as a preferred WWPN. If the controller has nothing specified, this is a blank field.
4. The *active_WWPN* indicates the WWPN of the storage controller port currently being used for I/O.
 - a. If no storage controller ports are available for I/O, this is a blank field.
 - b. If multiple controller ports are actively being used for I/O, this field's value is many.

online The MDisk is online and available.

(Internal MDisks only) The array has members that are degraded, or the raid status is degraded.

There are one or more MDisk port errors.

One or more paths to the MDisk have been lost; the MDisk is not online to every node in the system.

All paths to the MDisk are lost.

The MDisk is excluded from use by the system; the MDisk port error count exceeded the threshold.

```
lsmdisk -delim :
```

```
id:name:status:mode:mdisk_grp_id:mdisk_grp_name:capacity:ctrl_LUN_#:controller_name:UID:tier:encrypt:site_id:site_name:d
0:mdisk0:online:unmanaged:::100.0GB:0000000000000000:controller0:600a0b800076b42000002a1755e4f5e2000000000000000000000
1:mdisk1:degraded_paths:unmanaged:::1.0GB:0000000000000000:controller1:6005076802b580c10c000000000000000000000000000000
2:mdisk2:degraded_paths:managed:0:diskgrp2:1.0GB:0000000000000001:controller1:6005076802b580c10c000000000000020000000000
3:mdisk3:degraded_paths:unmanaged:::1.0GB:0000000000000002:controller1:6005076802b580c10c000000000000300000000000000000
4:mdisk4:degraded_paths:unmanaged:::1.0GB:0000000000000003:controller1:6005076802b580c10c000000000000400000000000000000
5:mdisk5:degraded_paths:unmanaged:::1.0GB:0000000000000004:controller1:6005076802b580c10c0000000000001000000000000000
```

```
lsmdisk mdisk1
```

```
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:600a0b80000f07020000005c45ff8a7c000000000000000000000000000000
preferred_WWPN:200400A0B80F0703
active_WWPN:200400A0B80F0703
fast_write_state:empty
raid_status:
raid_level:
redundancy:
strip_size:
spare_goal:
spare_protection_min:
balanced:
```

```
tier:tier0_flash
slow_write_priority:latency
fabric_type:fc
site_id:2
site_name:2
easy_tier_load:low
encrypt:no
distributed:no
drive_class_id
drive_count:8
stripe_width:4
total_rebuild_areas
available_rebuild_areas
rebuild_areas_goal
preferred_iscsi_port_id
active_iscsi_port_id
dedupe:yes
```


Ismdiskdumps (Deprecated)

The **lsdiskdumps** command is deprecated. Use the **lsdumps** command to display a list of files in a particular dumps directory.

Ismdisklba

Use the **lsmdisklba** command to list the MDisk and logical block address (LBA) for the specified volume LBA.

Syntax



Parameters

-vdisklba *vdisklba*

(Required) Specifies the 64-bit hexadecimal logical block address (LBA) on the volume. The LBA must be specified in hex, with a 0x prefix.

-copy *id*

(Optional) Specifies the volume copy ID to list the MDisk and LBA for. If this parameter is not specified, the command lists MDisk and LBAs for all volume 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 volume name or ID.

Description

The `lsmdisklba` command returns the logical block address (LBA) of the MDisk that is associated with the volume LBA. For mirrored volume, 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 volume and MDisk that are mapped in the same extent, or for thin-provisioned disks, in the same grain. If a thin-provisioned volume is offline and the specified LBA is not allocated, the command displays the volume LBA range only.

The `mdisk_lba` field provides the corresponding LBA on the real capacity for the input LBA. For compressed volume copies it is empty, and the system displays only the range of physical LBAs where the compressed input LBA is located.

Table 89 summarizes the data that can be returned with this command.

Table 89. `lsmdisklba` command output

Field	Fully allocated, single copy volume	LBA not allocated on thin-provisioned volume	Mirrored volume with one normal copy and one offline thin-provisioned copy	
			Normal copy	Thin-provisioned 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

`lsmdisklba -vdisk 0 -vdisklba 0x123`

The resulting output:

```
copy_id mdisk_id mdisk_name type      mdisk_lba      mdisk_start      mdisk_end      vdisk_start vdisk_end
0        1        mdisk1      allocated 0x0000000000100123 0x0000000000100000 0x00000000001FFFFF 0x00000000 0x000FFFFF
```

lsmdiskcandidate

Use the `lsmdiskcandidate` command to list all unmanaged MDisks by MDisk ID.

Syntax

➡ `lsmdiskcandidate` — `-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 storage pool or use a MDisk to create an image mode volume that the storage is actually used.

Check to see which MDisks are available by issuing the **detectmdisk** command to manually scan the Fibre Channel network for any MDisks. Issue the **lsmdiskcandidate** command to show the unmanaged MDisks. These MDisks have not been assigned to an storage pool. Alternatively, you can issue the **lsmdisk** command to view all of the MDisks.

An invocation example

```
lsmdiskcandidate
```

The resulting output:

```
id
5
6
7
8
9
10
11
12
13
14
```

lsmdiskextent

Use the **lsmdiskextent** command to display the extent allocation between managed disks and volumes. The output lists a volume ID, volume copy ID, and the number of extents.

Syntax

►► `lsmdiskextent` — `[-nohdr]` `[-delim delimiter]` `[mdisk_name | mdisk_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.

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 volume ID, volume copy ID, and the number of extents. These volume copies are using extents on the specified MDisk. The number of extents being used on each MDisk is also shown.

Every volume copy is constructed from one or more MDisk. 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 volume copies and MDisk, issue the following command for each volume copy:

```
lsdiskmember vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the volume copy. This displays a list of IDs that correspond to the MDisk that make up the volume copy.

To determine the relationship between volume copies and MDisk and the number of extents that are provided by each MDisk, you must use the command-line interface. For each volume copy, issue the following command:

```
lsdiskextent vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the volume copy. This displays a table of MDisk IDs and the corresponding number of extents that each MDisk is providing as storage for the given volume copy.

To determine the relationship between MDisk and volume copies, issue the following command for each MDisk:

```
lsmdiskmember 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 volume copies that are using this MDisk.

To determine the relationship between MDisks and volume copies and the number of extents that are used by each volume copy, you must use the command-line interface. For each MDisk, issue the following command:

```
lsmdiskextent mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This command displays a table of volume copy IDs and the corresponding number of extents that are being used by each volume copy.

An invocation example

```
lsmdiskextent -delim : mdisk0
```

The resulting output:

```
id:number_of_extents:copy_id  
1:1:1
```

lsmdiskmember

Use the **lsmdiskmember** command to display a list of volumes that use extents on the specified MDisk. That is, the volumes use extents on the managed disk that are specified by the MDisk ID.

Syntax

```
➤— lsmdiskmember — [ -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 name of the MDisk for which you want a list of volumes that use extents of that MDisk.

Description

This command displays a list of volumes 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 volume is constructed from one or more MDisks. To determine the relationship between volume copies and MDisks, issue the following command:

```
lsvdiskmember vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the volume copy. This displays a list of IDs that correspond to the MDisks that make up the volume copy.

To determine the relationship between volume copies and MDisks and the number of extents that are provided by each MDisk, you must use the command-line interface. For each volume copy, issue the following command:

```
lsvdiskextent vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the volume copy. This command displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the volume copy.

To determine the relationship between MDisks and volume copies, issue the following command:

```
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 volume copies that are using this MDisk.

To determine the relationship between MDisks and volume copies and the number of extents that are used by each volume copy, you must use the command-line interface. For each MDisk *mdisk_id* | *mdisk_name*, issue the following command:

```
lsmdiskextent mdisk_id | mdisk_name
```

where *mdisk_id* | *mdisk_name* is the name or ID of the MDisk. This command displays a table of volume copy IDs and the corresponding number of extents that are being used by each volume copy.

An invocation example

```
lsmdiskmember -delim : 1
```

The resulting output:

```
id:copy_id
0:0
1:0
2:0
3:0
4:0
5:0
6:0
```

setquorum (Deprecated)

The **setquorum** command is deprecated. Use the **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.

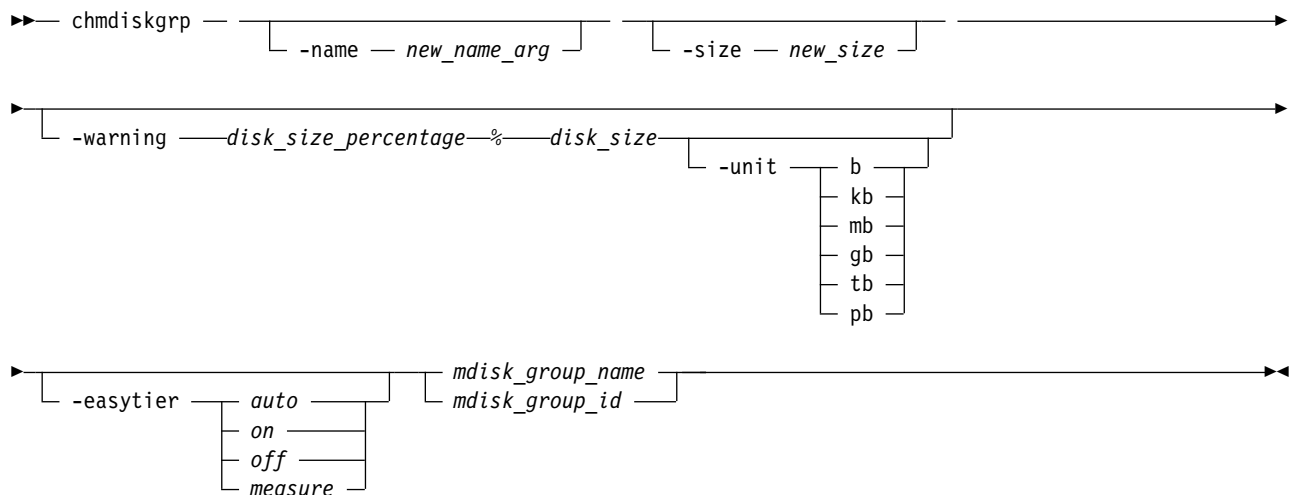
Chapter 21. Storage pool commands

Use the storage pool commands to work with storage pool options with the SAN Volume Controller.

chmdiskgrp

Use the **chmdiskgrp** command to modify the name that is assigned to a storage pool or to set the warning threshold for the storage pool.

Syntax



Parameters

-name *new_name_arg*
(Optional) Specifies the new name of the storage pool.

-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 storage pool. 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 storage pool size. To disable warnings, specify **0** or **0%**.

-size *new_size*
(Optional) Specifies the new size of a child pool.

Note: This parameter cannot be used with parent pools. Use **addmdisk** or **rmmdisk** to change storage pool capacity.

-unit **b** | **kb** | **mb** | **gb** | **tb** | **pb**
(Optional) Specifies the data units for the **-warning** parameter.

-easytier *auto* | *on* | *off* | *measure*
(Optional) Specifies if the Easy Tier function is on or off for this storage pool, or if it is automatically determined. **-easytier** is active in storage pools with multiple tiers and is balance with single tiers.

Note: **-easytier** must be followed by one of the following:

- If **-easytier** is set to *auto*, SAN Volume Controller automatically enables Easy Tier functions when the storage pool contains MDisks from more than one tier, and enables automatic rebalancing when the storage pool contains MDisks from only one tier.
- If **-easytier** is set to *on*, then Easy Tier functions are active.
- If **-easytier** is set to *off*, then Easy Tier functions are inactive.
- If **-easytier** is set to *measure* Easy Tier statistics are collected but Easy Tier management is disabled. (No extents are moved by Easy Tier.)

auto equates to:

- *on* if Easy Tier is licensed or no license is required
- *off* if Easy Tier is not licensed and a license is required

Specifying **-easytier on** enables Easy Tier:

- Management of both single-tier and multi-tier pools
- Auto rebalance

Extents are moved to balance the I/O load on the MDisks in the pool.

mdisk_group_id | *mdisk_group_name*

(Required) Specifies the ID or name of the storage pool to modify.

Description

Table 90. Parameter differences for child pools and parent pools

Parameter	Child pool usage	Storage pool usage
-name	Optional	Optional
-easytier	Cannot be used with child pools	Optional
-size	Optional	Cannot be used with parent pools
-unit	Optional	Optional
-warning	Optional	Optional

This command modifies the name, or label, assigned to a given storage pool. You can use the new name to refer to the storage pool.

The command can also be used to set the warning threshold for the storage pool. The warning threshold is the threshold at which a warning is generated when it is exceeded by the used-disk capacity in the storage pool.

You can also use this command to change other settings for parent pools and child pools.

An invocation example

```
chmdiskgrp -name testmdiskgrp -easytier on Group0
```

The resulting output:

No feedback

An invocation example

```
chmdiskgrp -size 100 -unit tb mypool
```

The resulting output:

No feedback

lsfreeextents

Use the **lsfreeextents** command to list the number of free extents that are available on a specified MDisk.

Syntax

```
➤— lsfreeextents — [ -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

```
lsfreeextents 2
```

The resulting output:


```
id 2
number_of_extents 4372
```

lsmdiskgrp

Use the **lsmdiskgrp** command to display a concise list or a detailed view of storage pools that are visible to the clustered system (system).

Syntax

```
➤— lsmdiskgrp — [ -filtervalue — attribute=value ] [ -nohdr ] [ -bytes ] ➤
```



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 when using the command-line interface (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:

```
lsmdiskgrp -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?

(Optional) Displays a list of valid filter attributes. The valid filters for the **lsmdiskgrp** command are:

- name
- id
- mdisk_count
- vdisk_count
- status
- storage_pool_id
- easy_tier
- easy_tier_status
- site_id

- site_name
- parent_mdisk_grp_id
- parent_mdisk_grp_name
- child_mdisk_grp_count
- type
- encrypt
- owner_type

Description

This command returns a concise list or a detailed view of storage pools visible to the system.

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 volume copies that are in the storage pool.

capacity

The total amount of MDisk storage that is assigned to the storage pool.

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 storage pool that is unused. MDisk storage can be used for system quorum data and volumes.

real_capacity

The amount of MDisk storage that is assigned to the storage pool that is assigned to volumes.

virtual_capacity

The total virtual size of all the volume copies that are associated with the storage pool. This is the same as the `real_capacity` value unless you have configured thin-provisioned volume copies in this storage pool.

used_capacity

The total used size of all the volume copies that are associated with the storage pool. This is the same as the `real_capacity` value unless you have configured thin-provisioned volume copies in this storage pool.

overalllocation

Expressed as a percentage, the ratio of the `virtual_capacity` value to the capacity. An storage pool overallocation of over 100 is only possible if you have configured thin-provisioned volume copies.

warning

This field is a percentage. A warning is generated when the amount of space in the storage pool 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: The values are:

1. on, which indicates Easy Tier actively manages the extents (including single-tier storage pools), and the Easy Tier status should be active - unless a license is required.

2. off, which indicates Easy Tier does not actively manage the extents, and the Easy Tier status should be inactive .
3. auto, which indicates the value of Easy Tier status is determined by the number of tiers an storage pool has.

Note: auto, which indicates:

- on if Easy Tier is licensed or no license is required
 - off if Easy Tier requires a license and none exists
4. measure indicates Easy Tier s collects statistics on that storage pool but does not move any extents in the storage pool.

easy_tier_status

Whether the Easy Tier functions are active on an storage pool:

- active indicates a pool is being managed by Easy Tier to provide tier management performance-based pool balancing (for example, extents can be moved).
- inactive indicates Easy Tier is inactive
- measured indicates Easy Tier statistics are being collected but no Easy Tier management is detected.
- balanced indicates a pool is being managed by Easy Tier to provide performance-based pool balancing (for example, extents can be moved).

This table describes the 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 523)
Off	One	On	inactive (see note 2 on page 523)
Off	Two	Off	inactive (see note 2 on page 523)
Off	Two	On	inactive (see note 2 on page 523)
Measure	One	Off	measured (see note 3 on page 523)
Measure	One	On	measured (see note 3 on page 523)
Measure	Two	Off	measured (see note 3 on page 523)
Measure	Two	On	measured (see note 3 on page 523)
Auto	One	Off	measured (see note 3 on page 523)
Auto	One	On	balanced (see note 4 on page 523)
Auto	Two	Off	measured (see note 3 on page 523)
Auto	Two	On	active (see note 5 on page 523)
On	One	Off	balanced (see note 4 on page 523)

Storage pool Easy Tier setting	Number of tiers in the storage pool	Volume copy Easy Tier setting	Volume copy Easy Tier status
On	One	On	balanced (see note 4)
On	Two	Off	measured (see note 3)
On	Two	On	active (see note 5)

Note:

1. If the volume copy is in image or sequential mode or is being migrated then the volume copy Easy Tier(tm) status is measured instead of active.
2. When the volume copy status is inactive, no Easy Tier(tm) functions are enabled for that volume copy.
3. When the volume copy status is measured, the Easy Tier(tm) function collects usage statistics for the volume but automatic data placement is not active.
4. When the volume copy status is balanced, the Easy Tier(tm) function enables performance-based pool balancing for that volume copy.
5. When the volume copy status is active, the Easy Tier(tm) function operates in automatic data placement mode for that volume.
6. 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 other than pool performance balancing are disabled for storage pools with a single tier, and that automatic data placement mode is enabled for all striped volume copies in a storage pool with two or more tiers.

tier Indicates which tier information is being reported. The values are:

- tier0_flash
- tier1_flash
- tier_enterprise
- tier_nearline

tier_mdisk_count

Indicates the number of MDisk in the tier.

tier_capacity

Indicates the total MDisk capacity assigned to the volume in the tier.

Note: For thin-provisioned copies, the capacity by tier will be the real capacity.

tier_free_capacity

Indicates the unused amount of MDisk storage in the tier.

compression_active

Indicates whether there are any compressed volume copies in the storage pool.

compression_virtual_capacity

Indicates the total virtual capacity for all compressed volume copies in the storage pool. This is in unsigned decimal format.

compression_compressed_capacity

Indicates the total used capacity for all compressed volume copies in the storage pool. This is in unsigned decimal format.

compression_uncompressed_capacity

Indicates the total uncompressed used capacity for all compressed volume copies in the storage pool. This is in unsigned decimal format.

site_id

Indicates the site value for the storage pool group. This numeric value is 1, 2, 3 or blank.

site_name

Indicates the site name for the storage pool. This is an alphanumeric value or is blank.

parent_mdisk_grp_id

Indicates the storage pool group ID. This is a numeric string (between 0 and 127 characters) or blank.

parent_mdisk_grp_name

Indicates the storage pool group name. This is an alphanumeric string (between 1 and 63 characters) or blank.

child_mdisk_grp_count

Indicates the number of child pools in the parent pools. This is a numeric string (between 0 and 127 characters) or blank.

child_mdisk_grp_capacity

Indicates the total amount of space reserved for child pools.

type Indicates the MDisk group type. The values are parent and child_thick.

encrypt

Indicates whether the data stored on the MDisk group is encrypted or not encrypted. The values are:

- yes if the pool has an encryption key
- yes if the pool does not have an encryption key (and the pool contains MDisk and all are encrypted)
- no if the pool does not have an encryption key (and the pool contains MDisk and at least one is not encrypted)
- Blank if the pool does not have an encryption key (and the pool has no MDisk)

The following define the status fields, from lowest to highest priority:

Online

Indicates that the storage pool is online and available.

Offline

Indicates that all paths to the storage pool are lost.

owner_type

Indicates the type of owning object, such as a file system or application. This is an alphanumeric string up to 20 characters in length.

owner_id

Indicates an identifier for the owning object. It is represented by a number and is blank if there is no object

owner_name

Indicates the name for the object that owns the volume. This is an alphanumeric string up to 63 characters in length or is blank.

A concise invocation example

```
lsmdiskgrp -delim :
```

The concise resulting output:

```
id:name:status:mdisk_count:vdisk_count:capacity:extent_size:free_capacity:virtual_capacity:used_capacity:real_capacity:overallocation:0:mdiskgrp0:online:2:1:42.18GB:512:8.00GB:34.18GB:34.18GB:34.18GB:81:0:measure:measured:no:0.00MB:0.00MB:0.00MB:yes:1:spnrgfld1
```

A detailed invocation example for a storage pool with one tier

```
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:ssd
tier_mdisk_count:0
tier_capacity: 0.00MB
tier_free_capacity:0.00MB
tier_tier0_flash
tier_mdisk_count 1
tier_capacity 1.63TB
tier_free_capacity 1.63TB
tier_tier1_flash
tier_mdisk_count 1
tier_capacity 1.63TB
tier_free_capacity 1.63TB
tier_tier_enterprise
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB
tier_tier_nearline
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB
compression_active:yes
compression_virtual_capacity:1000.00MB
compression_compressed_capacity:0.41MB
compression_uncompressed_capacity:512.05MB
site_id:3
site_name:Quorum
parent_mdisk_grp_id:3
parent_mdisk_grp_name:sisfyle
child_mdisk_grp_count:0
child_mdisk_grp_capacity:0.00MB
type:child_thick
encrypt:no
owner_type vvol_child_pool
owner_id
owner_name

```

A detailed invocation example for a storage pool with two tiers

```
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:ssd
tier_mdisk_count:2
tier_capacity:20.00GB
tier_free_capacity:0.00MB
tier_tier0_flash
tier_mdisk_count 1
tier_capacity 1.63TB
tier_free_capacity 1.63TB
tier_tier1_flash
tier_mdisk_count 1
tier_capacity 1.63TB
tier_free_capacity 1.63TB
tier_tier_enterprise
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB
tier_tier_nearline
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB

tier_mdisk_count:6
tier_capacity:180.00GB
tier_free_capacity:100.00GB
tier:ri_ssd
tier_mdisk_count:
tier_capacity:
tier_free_capacity:
compression_active:yes
compression_virtual_capacity:1000.00MB
compression_compressed_capacity:0.41MB
compression_uncompressed_capacity:512.05MB
site_id:2
site_name:POK
parent_mdisk_grp_id:2
parent_mdisk_grp_name:sysfile
child_mdisk_grp_count:0
child_mdisk_grp_capacity:0.00MB
type:child_thick
owner_type vvol_child_pool
owner_id
owner_name

```

A detailed invocation example for a storage pool with three tiers

```
lsmdiskgrp -delim : mdiskgrp1
```

The resulting output:

```

id:1
name:mdiskgrp1
status:online
mdisk_count:4
vdisk_count:6
capacity:200.00GB
extent_size:16
free_capacity:100.00GB
virtual_capacity:400.00GB
used_capacity:75.00GB
real_capacity:100.00GB
overallocation:200
warning:80
easy_tier:auto
easy_tier_status:inactive

```



```

tier tier0_flash
tier_mdisk_count 1
tier_capacity 1.63TB
tier_free_capacity 1.63TB
tier tier1_flash
tier_mdisk_count 1
tier_capacity 1.63TB
tier_free_capacity 1.63TB
tier tier_enterprise
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB
tier tier_nearline
tier_mdisk_count 0
tier_capacity 0.00MB
tier_free_capacity 0.00MB

compression_active:no
compression_virtual_capacity:0.00MB
compression_compressed_capacity:0.00MB
compression_uncompressed_capacity:0.00MB
site_id:2
site_name:POK
parent_mdisk_grp_id:1
parent_mdisk_grp_name:filesys
child_mdisk_grp_count:0
child_mdisk_grp_capacity:0.00MB
type:child_thick

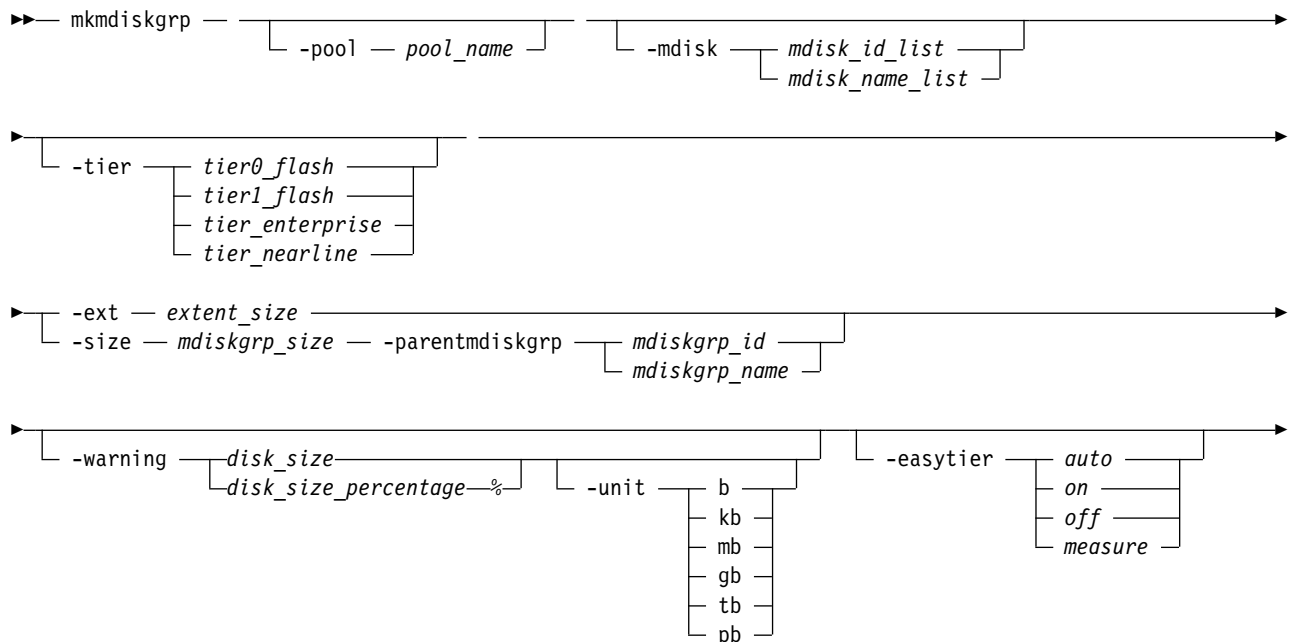
owner_type vvol_child_pool
owner_id
owner_name


```

mkmdiskgrp

Use the **mkmdiskgrp** command to create a new storage pool.

Syntax





Parameters

- name** *pool_name*
(Optional) Specifies a name to assign to the new pool.
- mdisk** *mdisk_id_list* | *mdisk_name_list*
(Optional) Specifies a colon-separated list of managed disk IDs or names to add to the storage pool. You can create an empty storage pool by not specifying the **-mdisk** parameter.
- tier** *tier0_flash* | *tier1_flash* | *tier_enterprise* | *tier_nearline*
(Optional) Specifies the tier of the MDisk or MDisk being added.
 - tier0_flash**
Specifies a tier0_flash hard disk drive or an external MDisk for the newly discovered or external volume.
 - tier1_flash**
Specifies an tier1_flash (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.
 - tier_enterprise**
Specifies a tier_enterprise hard disk drive or an external MDisk for the newly discovered or external volume.
 - tier_nearline**
Specifies a tier_nearline hard disk drive or an external MDisk for the newly discovered or external volume.

If you do not specify a tier, the current tier value of the MDisk is retained. The default value for an external MDisk is *enterprise*.
- 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).
- size** *mdiskgrp_size*
(Optional) Specifies the child pool capacity. The value must be a numeric value (and an integer multiple of the extent size).
- parentmdiskgrp** *mdiskgrp_id* | *mdiskgrp_name*
(Optional) Specifies the parent pool from which the volume extents of the child pool are allocated when creating a child pool. The value must be an *mdiskgrp_id* or *mdiskgrp_name*.
- warning** *disk_size* | *disk_size_percentage*
(Optional) Generates a warning when the used disk capacity in the storage pool 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 storage pool 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.
- easytier** *on* | *off* | *auto* | *measure*
(Optional) Specifies if the Easy Tier function is active for this storage pool, or if it is automatically determined. *auto* is the default value. **-easytier** is active in storage pools with multiple tiers and is balance with single tiers.

Note:

- If **-easytier** is set to *auto*, the system automatically enables Easy Tier functions when the storage pool contains MDisk from more than one tier, and enables automatic rebalancing when the storage pool contains an MDisk from only one tier.
- If **-easytier** is set to *on*, then Easy Tier functions are active.
- If **-easytier** is set to *off*, then Easy Tier functions are inactive.
- If **-easytier** is set to *measure* Easy Tier statistics are collected but Easy Tier management is disabled. (No extents are moved by Easy Tier.)

auto equates to:

- *on* if Easy Tier is licensed or no license is required
- *off* if Easy Tier is not licensed and a license is required

Specifying **-easytier on** enables Easy Tier:

- Management of both single-tier and multi-tier pools
- Auto rebalance

Extents are moved to balance the I/O load on the MDisks in the pool.

-owner *owner_type*

(Optional) Specifies the owner type. The value must be *vvol_child_pool*.

-encrypt *yes | no*

(Optional) Specifies the encryption status for this storage pool. The values are *yes* or *no*.

Remember:

- If you do not specify **-encrypt** and encryption is enabled, the system defaults to **-encrypt yes** (the default encryption setting).
- When creating a child pool in an encrypted parent pool, the value must not be *no*. (All other permutations are permitted.)

Description

Table 91. Parameter differences for child pools and storage pools

Parameter	Child pool usage	Storage pool usage
-name	Optional	Optional
-mdisk	Cannot be used with child pools	Optional
-tier	Cannot be used with child pools.	Optional
-easytier	Cannot be used with child pools	Optional
-size	Mandatory	Cannot be used with parent pools
-parentmdiskgrp	Mandatory	Cannot be used with parent pools
-ext	Cannot be used for child pools	Mandatory
-unit	Optional	Optional
-warning	Optional	Optional
-encrypt	Optional	Optional for both parent pools and child pools

The **mkmdiskgrp** command creates a new storage pool and assigns the storage pool name if specified. The ID of the new storage pool is returned if the command is successful. Storage pools are collections of managed disks. Each storage pool is divided into chunks, called extents, which are used to create volumes.

Optionally, you can specify a list of managed disks that will be added to this storage pool. These managed disks cannot belong to another storage pool, and they must have a mode of unmanaged. Use the **lsmdiskcandidate** 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 volume is created from this group, free extents from the pool are used, in accordance with the policy used when the volume 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 volumes and managed disks. The clustered system (system) can only manage a finite number of extents (4 194 304). One system can virtualize the following number of extents:

- 64 TB – if all storage pools have extent sizes of 16 MB.
- 2 PB – if all storage pools have extent sizes of 512 MB.
- 32 PB – if all storage pools have extent sizes of 8192 MB.

Important: The extent size for the storage pool can also limit volume size. Consider the maximum volume size you want to use when creating storage pools. Refer to the information on creating storage pools for a comparison of the maximum volume capacity for each extent size. The maximum is different for thin-provisioned volumes.

Note: When an image mode volume is created, the storage pool increases in capacity by the size of the image mode volume (not the MDisk capacity), because the image mode volume might be smaller than the MDisk itself. If an extent is migrated from the image mode volume or MDisk to elsewhere in the group, the volume becomes a striped volume (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 volume) becomes available.

When specifying **-name pool_name** if you do not also specify **-parentmdiskgrp**, you create a parent pool where *pool_name* is the name of the new storage pool. When specifying **-name pool_name**, if you also specify **-parentmdiskgrp** and a size for it, you create a child pool where *pool_name* is the name of the new storage pool.

An invocation example

This example adds a list of MDisks to the storage pool.

```
mkmdiskgrp -mdisk mdisk0:mdisk1:mdisk2:mdisk3 -ext 32
```

The resulting output:

```
MDisk Group, id [0], successfully created
```

An invocation example

This example specifies tier and Easy Tier information when adding a list of MDisks to the storage pool.

```
mkmdiskgrp -mdisk mdisk13:mdisk16 -ext 512 -tier tier_nearline -easytier measure
```

The resulting output:

```
MDisk Group, id [13], successfully created
```

An invocation example

This example creates a child pool from a parent pool.

```
mkmdiskgrp -size 100 -unit tb -parentmdiskgrp phyppool
```

The resulting output:

```
MDisk Group, id [3], successfully created
```

An invocation example

This example creates a child pool from a parent pool and specifies an owner type.

```
mkmdiskgrp -parentmdiskgrp p0 -size 100 -unit gb -owner vvol_child_pool
```

The resulting output:

```
MDisk Group, id [3], successfully created
```

An invocation example

This example creates an encrypted child pool from a parent pool.

```
mkmdiskgrp -parentmdiskgrp 2 -name _my_encrypted_child_pool -encrypt yes -size 10 -unit gb
```

The resulting output:

```
MDisk Group, id [5], successfully created
```

rmmdisk

Use the **rmmdisk** command to delete a managed disk (MDisk) from a storage pool.

Syntax

```
➤— rmmdisk — — -mdisk ———— mdisk_id_list ———— ———— mdisk_group_id ———— ➤  
                        └── mdisk_name_list ┘                └── -force ┘                └── 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 storage pool to delete the disks from. The warning threshold for a storage pool is automatically scaled when MDiskS are deleted.

Description

This command attempts to remove the managed disk or disks from the group.

Remember: This command cannot be used for child pools.

Deleting a managed disk from a group can only be done if the managed disk does not contain any extents in use by a volume. 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 MDisk 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 storage pool that they belong to.

The **rmmdisk** command fails if there are insufficient free extents on other disks in the storage pool 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 storage pool. If there are not enough free extents in the storage pool, 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 volume 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 (if migration is required). 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 **lsmigrate** command.

If you specify **-force**, the **rmmdisk** command fails if there are offline MDisk. If there are no online quorum disks the migration fails.

Remember: When using the **-mdisk** parameter, MDisk are removed if there is one (or more) SAS MDisk specified in the list.

An invocation example

```
rmmdisk -mdisk mdisk12 -force Group3
```

The resulting output:

No feedback

rmmdiskgrp

Use the **rmmdiskgrp** command to delete a storage pool without being able to recover it.

Syntax

```
➤— rmmdiskgrp — [ -force ] [ mdisk_group_id | mdisk_group_name ] —➤
```

Parameters

-force

(Optional) Specifies that all volumes and host mappings be deleted. When you use this parameter, all managed disks in the storage pool are removed and the storage pool itself is deleted.

Remember:

- You must specify **-force** to delete a child pool if it contains volume.
- You cannot specify **-force** to delete a parent pool if it has child pools.

mdisk_group_id | mdisk_group_name

(Required) Specifies the ID or name of the storage pool that is to be deleted.

Note: You cannot delete a parent pool that has child pools. You must first delete the child pools.

Description

Important: Before you issue the command, ensure that you want to delete all mapping information; data that is contained on volume cannot be recovered after the storage pool has been deleted

The **rmmdiskgrp** command deletes the specified storage pool. The **-force** parameter is required if there are volumes that have been created from this storage pool or if there are managed disks in the storage pool. Otherwise, the command fails.

Note: This command also removes any associated storage pool throttling.

Deleting a storage pool is essentially the same as deleting a clustered system (system) or part of a system, because the storage pool is the central point of control of virtualization. Because volumes are created using available extents in the storage pool, mapping between volume extents and managed disk extents is controlled based on the storage pool.

The command deletes all volume copies in the specified storage pool. If the volume has no remaining synchronized copies in other storage pools, the volume is also deleted.

This command deletes the associated MDisk group (storage pool) throttle if that storage pool is removed.

Remember: This command is unsuccessful if:

- Volume protection is enabled (using the **chssystem** command)
- The MDisk being removed is mapped to any volume that has received I/O within the defined volume protection time period

Remember: This command partially completes asynchronously. All volumes, host mappings, and Copy Services relationships are deleted before the command completes. The deletion of the storage pool then completes asynchronously.

In detail, if you specify the **-force** parameter and the volumes are still using extents in this storage pool, 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 volume 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 system to also be stopped. See the description for the **-force** parameter in the **stopfcmap** command for additional information.

- Any volume that is in the process of being migrated into or out of the storage pool is deleted. This frees up any extents that the volume was using in another storage pool.
- Volumes are deleted without first flushing the cache. Therefore, the storage controller LUNs that underlie any image mode MDisk might not contain the same data as the image mode volume prior to the deletion.
- If there are managed disks in the storage pool, all disks are deleted from the storage pool. They are returned to the unmanaged state.
- The storage pool is deleted.

Attention: If you use the **-force** parameter to delete all the storage pools in your system, you are returned to the processing state where you were after you added nodes to the system. All data that is contained on the volumes is lost and cannot be recovered.

An invocation example

```
rmmdiskgrp -force Group3
```

The resulting output:

No feedback

Chapter 22. Copy services commands

The Copy Service commands enable you to work with the Metro Mirror, Global Mirror, and HyperSwap services that the SAN Volume Controller provides.

chpartnership

Use the **chpartnership** command to modify the bandwidth of the partnership between the local clustered system (system) and the remote system that is specified in the command. This affects the bandwidth that is available for background copy in a system partnership by either Metro Mirror or Global Mirror operations. Additionally, use this command to disable and re-enable the partnership, which allows the local system to be disconnected and then reconnected to the remote system.

Syntax

```
►► chpartnership — [ -start ] [ -stop ] [ remote_cluster_id ] [ remote_cluster_name ]
[ -linkbandwidthmbits — link_bandwidth_in_mbps ]
[ remote_cluster_id ] [ remote_cluster_name ]
►► chpartnership — [ -type [ ipv4 | ipv6 ] ] [ -clusterip [ newipv4addr | newipv6addr ] ]
[ -chapsecret — newCHAPsecret ] [ -nochapsecret ]
[ -backgroundcopyrate — percentage ] [ -linkbandwidthmbits — link_bandwidth_in_mbps ]
[ -compressed [ yes | no ] ] [ remote_cluster_id ] [ remote_cluster_name ]
```

Parameters

-start | -stop

(Optional) Starts or stops a Metro Mirror or Global Mirror partnership. To start or stop a partnership, run the **chpartnership** command from either system.

-type *ipv4* | *ipv6*

(Optional) Specifies the Internet Protocol (IP) address format for the partnership using either of these case-sensitive strings:

- *ipv4* for Internet Protocol Version 4 (IPv4)
- *ipv6* for Internet Protocol Version 6 (IPv6)

This migrates a partnership from *ipv4* to *ipv6* or vice versa.

-clusterip *newipv4addr | newipv6addr*

(Optional) Specifies the new partner system IP address, either *ipv4* or *ipv6*. Systems connected over IP links are not displayed by **lspartnershipcandidate** before executing **mkippartnership**. This does not apply to FC-based or FCoE-based connections.

Specify this parameter when creating partnerships with systems connected over native IP links. To change the partner system IP address, first specify **chpartnership -stop** to stop the partnership.

-chapsecret *newCHAPsecret*

(Optional) Specifies the new Challenge-Handshake Authentication Protocol (CHAP) secret of the partner system. The maximum size of the CHAP secret is eighty alphanumeric characters.

-nochapsecret

(Optional) Resets the CHAP secret used to authenticate with the partner system. Specify **chpartnership -stop** to stop the partnership. Reset the CHAP secret of the partner system when authentication of discovery requests is turned off on the partner system (by specifying **chsystem -rcauthmethod**).

-backgroundcopyrate *percentage*

(Optional) Specifies the maximum percentage of aggregate link bandwidth that can be used for background copy operations. This parameter can be specified without stopping the partnership. The percentage is a numeric value from 0 to 100, and the default value is 50, which means that a maximum of 50% of the aggregate link bandwidth can be used for background copy operations. This command is mutually-exclusive with all parameters other than **-linkbandwidthmbits**.

Note: If the specified value is non-zero, the combination of both the **-backgroundcopyrate** and the **-linkbandwidthmbits** values must result in a background copy bandwidth of at least 8 megabits per second (Mbps).

-linkbandwidthmbits *link_bandwidth_in_mbps*

(Optional) Specifies the aggregate bandwidth of the RC link between two clustered systems (systems) in megabits per second (Mbps). It is a numeric value from 1 to 100000, specified in Mbps. This parameter can be specified without stopping the partnership. This command is mutually-exclusive with all parameters other than **-backgroundcopyrate**.

Note: If the specified value is non-zero, the combination of both the **-backgroundcopyrate** and the **-linkbandwidthmbits** values must result in a background copy bandwidth of at least 8 Mbps.

remote_cluster_ID | remote_cluster_name

(Required) Specifies the remote system ID or name of a partnership. The specified value must match one of the system IDs or names returned after issuing **lspartnershipcandidate**. The specified value must match one of the system IDs or names listed by **lspartnership**.

Remember: Specifying a remote system ID or name with **chpartnership** does not affect the remote system. To change the system name, specify **chsystem**.

To configure the maximum bandwidth available for Metro Mirror intrasystem relationships, specify:

- A local system ID or name
- The **-linkbandwidthmbits** and **-backgroundcopyrate** parameters

-compressed *yes | no*

(Optional) Specifies whether compression is enabled for this partnership. The default value is no.

Description

This command modifies the bandwidth of the partnership between the local system and the remote system specified in the command. This affects the bandwidth available for a background copy in Metro Mirror or Global Mirror relationships (from the local to the remote system) . To modify the background copy bandwidth from remote system to local system, issue **chpartnership** a second time for the remote system.

Change the CHAP secret or system IP for partnerships created using IP links. Before changing the partner CHAP secret or system IP, stop the partnership.

If a stop partnership is issued, the state is `not_present` briefly before it changes to `fully_configured_stopped`.

Important:

- If you start with a fully configured remote copy partnership, the state (as reported by `lspartnership`) is `fully_configured`.
- If a stop partnership is issued, the state is `not_present` (typically for ten seconds or less) before it becomes `fully_configured_stopped`.

After making the necessary changes, start the partnership.

The system partnership must be in either the `partially_configured_stopped` or `fully_configured_stopped` states to be started.

Note: The local and remote systems in an IP partnership must use the same IP address types, IPv4 or IPv6.

An invocation example

```
chpartnership -stop cluster1
```

The resulting output:

No feedback

An invocation example to change the allocated background copy rate

```
chpartnership -backgroundcopyrate 20 remote-system-2
```

The resulting output:

No feedback

An invocation example to change the link bandwidth

```
chpartnership -linkbandwidthmbits 1024 remote-system-2
```

The resulting output:

No feedback

An invocation example to migrate existing partnership from IPv4 to IPv6 type

```
chpartnership -stop remote-sys-2
chpartnership -type ipv6 -clusterip fe80::200:f8ff:fe21:67cf remote-sys-2
```

The resulting output:

No feedback

An invocation example to configure a new CHAP secret for a partner

```
chpartnership -stop remote-system-2
chpartnership -chapsecret newpassword remote-system-2
chpartnership -start remote-system-2
```

The resulting output:

No feedback

An invocation example to configure a new system IP

```
chpartnership -stop remote-system-2
chpartnership -clusterip 202.49.86.2 -chapsecret newpassword remote-system-2
chpartnership -start remote-system-2
```

The resulting output:

No feedback

An invocation example setting the aggregate bandwidth and background copy rate

```
chpartnership -linkbandwidthmbits 2048 -backgroundcopyrate 100 localCluster
```

The resulting output:

No feedback

An invocation example enabling compression on an IP replication link

```
chpartnership -compressed yes svtcluster1
```


The resulting output:

No feedback

chrconsistgrp

Use the **chrconsistgrp** command to modify attributes of an existing Metro Mirror, Global Mirror, or active-active consistency group, such as changing the name of a consistency group.

Syntax



Parameters

-name *new_name_arg*

(Optional) Specifies the new name to assign to the consistency group.

-cycleperiodseconds *period*

(Optional) Specifies the cycle period in seconds. The minimum cycle period value is 60 seconds, and the default is 300 seconds.

This defines an optional cycle period that applies to Global Mirror relationships with a cycling mode of *multi*. A Global Mirror relationship using the *multi* cycling_mode performs a complete cycle each period. It might be provided for any relationship, but cannot be used for *none* when considering Metro or Global Mirror relationships.

-cyclingmode *none* | *multi*

(Optional) Specifies the behavior of Global Mirror for this relationship.

- Specifying *none*, the default, gives identical behavior to Global Mirror in previous versions of SAN Volume Controller.
- Specifying *multi* uses the cycling protocol.

To start a relationship with `cycling_mode` set to *multi*, change volumes must be defined for the relationship.

Note: The `cycling_mode` can only be changed when the relationship is stopped and in `consistent_stopped` or `inconsistent_stopped` states.

-metro

(Optional) Specifies the change in the consistency group's copy type and converts a Global Mirror (with or without change volumes) relationship to a Metro Mirror relationship.

Remember: To use this parameter the consistency group must be stopped (`inconsistent_stopped`, `consistent_stopped`, or `idling`)

-global

(Optional) Specifies the change in the consistency group's copy type and converts a Metro Mirror relationship to a Global Mirror relationship. This parameter is not mutually-exclusive with **-cyclingmode**. If you do not specify **-cyclingmode** and the relationship is Metro Mirror, the `cycling_mode` value is none.

Remember: To use this parameter the consistency group must be stopped (`inconsistent_stopped`, `consistent_stopped`, or `idling`)

`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 modifies the specified attributes of the supplied consistency group, one attribute at a time.

All parameters are mutually-exclusive with the exception of the **-cyclingmode**, which is mutually-exclusive with all parameters but **-global**.

Note: One of the optional parameters must be specified.

You can change a relationship or consistency group between copy types even if replication is stopped. Consistency protection is preserved across all types, so a relationship or consistency group that is in `consistent_copying` state before being stopped retains the consistent copy on the secondary system when the copying type is changed.

Note: You cannot set cycling mode to multi-cycling mode if there is a relationship where the primary and secondary volume are different sizes.

A Global Mirror consistency group with cycling mode set to *multi* requires that change volumes are defined for the primary and secondary volumes of each relationship in the group before it can be started.

For intersystem relationships the;

- **-cycleperiodseconds** and **-cyclingmode** parameters can only be specified when the two systems are connected. If the two systems become disconnected while the command is being processed, then the command might be completed with the change having been performed at the system that received the task invocation only (and the other system is updated upon re-connection).

For consistency groups that are active-active, you cannot change the copy type or cycling mode. This means you cannot specify these parameters:

- **-global**
- **-metro**
- **-cyclingmode**

```
chrconsistgrp -name rctestone rc testgrp
```

No feedback

Use the **chrcrelationship** command 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.

[illegible]

```
-masterchange master change vdisk id | master change vdisk name
```

(Optional) Specifies a change volume association for the master volume in the relationship.

```
-auxchange aux_change_vdisk_id | aux_change_vdisk_name
```

(Optional) Specifies a change volume association for the auxiliary volume in the relationship.

-nomasterchange

(Optional) Specifies a defined change volume on the master volume should be removed from the relationship.

Note: To use this parameter the specified change volume must no longer be in use by the relationship, including change volumes of a running relationship (`inconsistent_copying`, `consistent_copying`, or `consistent_synchronized`).

This does not include a primary change volume of a stopped relationship. A secondary change volume of a relationship stopped from `consistent_copying` is considered in use if the change volume is supplying the consistent image. If this change volume needs to be removed, the relationships must first be stopped using the **-access** parameter in order to apply the consistent image to the secondary volume.

-noauxchange

(Optional) Specifies a defined change volume on the auxiliary volume should be removed from the relationship.

Note: To use this parameter the specified change volume must no longer be in use by the relationship, including change volumes of a running relationship (*inconsistent_copying*, *consistent_copying*, or *consistent_synchronized*).

This does not include a primary change volume of a stopped relationship. A secondary change volume of a relationship stopped from *consistent_copying* is considered in use if the change volume is supplying the consistent image. If this change volume needs to be removed, the relationships must first be stopped using the **-access** parameter in order to apply the consistent image to the secondary volume.

-name *new_name_arg*

(Optional) Specifies a new label to assign to the relationship.

-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 (Global Mirror, Metro Mirror, or active-active) can be assigned to the same consistency group.

-noconsistgrp

(Optional) Removes the specified relationship from a consistency group, making the relationship a standalone relationship.

-cycleperiodseconds *period*

(Optional) Specifies the cycle period in seconds. The minimum cycle period value is 60 seconds. The default is 300 seconds (5 minutes).

This defines an optional cycle period that applies to Global Mirror relationships with a cycling mode of *multi*. A Global Mirror relationship using the *multi* *cycling_mode* performs a complete cycle at most once each period.

-cyclingmode *none* | *multi*

(Optional) Specifies the behavior of Global Mirror for this relationship.

- Specifying *none*, the default, gives identical behavior to Global Mirror in previous versions of SAN Volume Controller.
- Specifying *multi* uses the cycling protocol.

To start a relationship with *cycling_mode* set to *multi*, change volumes must be defined for the relationship.

Note: The *cycling_mode* can only be changed when the relationship is stopped and in *consistent_stopped* or *inconsistent_stopped* status.

-metro

(Optional) Specifies the change in the relationship's copy type and converts a Global Mirror (with or without change volumes) relationship to a Metro Mirror relationship.

Remember: To use this parameter the relationship must be stopped (*inconsistent_stopped*, *consistent_stopped*, or *idling*)

-global

(Optional) Specifies the change in the relationship's copy type and converts a Metro Mirror relationship to a Global Mirror relationship. This parameter is not mutually-exclusive with **-cyclingmode**. If you do not specify **-cyclingmode** and the relationship is Metro Mirror, the *cycling_mode* value is *none*.

Remember: To use this parameter the relationship must be stopped (`inconsistent_stopped`, `consistent_stopped`, or `idling`)

`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. In addition to changing the name of a consistency group, this command can be used for the following purposes.

Remember:

- All parameters are mutually-exclusive with the exception of the **-cyclingmode**, which is mutually-exclusive with all parameters but **-global**.
- One of the optional parameters must be specified.

Note: You cannot set cycling mode to multi-cycling mode if the primary and secondary volumes are different sizes.

- 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 system
 - Auxiliary system
 - State (unless the group is empty)
 - Primary (unless the group is empty)
 - Type (unless the group is empty)
 - Cycling mode (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), type (Metro Mirror or Global Mirror), and cycling mode 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 **-noconsistgrp** 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 systems are disconnected the relationship is only removed from the consistency group on the local system, at the time the command is issued. When the systems are reconnected the relationship is automatically removed from the consistency group on the other system. Alternatively, you can issue an explicit modify (**chrcrelationship**) command to remove the relationship from the group on the other system 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 **chrcrelationship** command twice. Use the **-noconsistgrp** parameter to remove the relationship from its current group, and then use the **-consistgrp** parameter with the name of the new consistency group.

You can change a relationship or consistency group between copy types even if replication is stopped. Consistency protection is preserved across all types, so a relationship or consistency group that is in `consistent_copying` state before being stopped retains the consistent copy on the secondary system when the copying type is changed.

If you add a `consistent_stopped` relationship that uses consistency protection to a consistency group that is not using consistency protection, the system attempts to trigger consistency protection on the consistency group that is not currently using consistency protection. If the relationship or consistency group that is not currently using consistency protection does not have secondary change volumes that are defined, specifying `chcrelationship -consistgrp` fails. If you add a relationship to a consistency group where at least one is using consistency protection, the resulting consistency group is not mutually consistent, which means that the data on the secondary system for the relationship being added is inconsistent with the data on the consistency group. This also means that enabling access to the volume fails.

For intersystem relationships

- The **-name**, **-consistgrp**, **-cycleperiodseconds** and **-cyclingmode** parameters can only be specified when the two systems are connected. If the two systems become disconnected while the command is being processed, then the command might be completed with the change having been performed at the system that received the task invocation only (and the other system is updated upon re-connection). The **-cycleperiodseconds** and **-cyclingmode** parameters can only be specified on stand-alone relationships (not members of a consistency group).
- The **-masterchange** and **-nomasterchange** parameters can only be specified when running the **chcrelationship** command on the master system for the relationship, and the **-auxchange** and **-noauxchange** parameters can only be specified when running the **chcrelationship** command on the auxiliary system for the relationship.

Remember: You cannot specify a master and auxiliary change volume in the same command.

A change volume must be:

- Used by the relationship that owns it
- In the same I/O group as the associated master or auxiliary volume
- The same size as the associated master or auxiliary volume

A change volume is owned and used by the associated Remote Copy relationship. Consequently, it cannot be:

- Mapped to a host
- Used as source or target of any FlashCopy maps
- Part of any other relationship
- A filesystem disk

Assigning a change volume to a relationship requires new FlashCopy mappings to be created between the master or auxiliary volume and the associated change volume. Consequently, there must be sufficient unallocated FlashCopy memory in the target I/O group or the command fails.

Note: You cannot use this command if cloud snapshot is enabled on the volume or the volume owner type is `cloud_backup`.

If the `cycle_period_seconds` for the relationship does not match that of the consistency group it is added to, the newly-added relationship copies the `cycle_period_seconds` value from the group. If later removed from the group, the copied `cycle_period_seconds` value remains.

When a Global Mirror relationship with a `cycling_mode` value of *multi* is added to a group that is not empty, both the group and the relationship must be stopped.

For relationships that are active-active, you cannot change the copy type, cycling mode, or the change volumes. This means you cannot specify these parameters:

- **-global**
- **-metro**

- **-cyclingmode**
- **-nomasterchange**
- **-noauxchange**

Remember: For a volume to be configured as a change volume for an active-active relationship, the volume must have the same site name and site ID as the master and auxiliary volume it is being associated with.

An invocation example to change a relationship name from rccopy1 to testrel

```
chrcrelationship -name testrel rccopy1
```

The resulting output:

No feedback

An invocation example to add relationship rccopy2 to group newgroup

```
chrcrelationship -consistgrp newgroup rccopy2
```

The resulting output:

No feedback

An invocation example to remove relationship rccopy3 from whichever consistency group it is a member of

```
chrcrelationship -noconsistgrp rccopy3
```

The resulting output:

No feedback

An invocation example

```
chrcrelationship -cyclingmode multi relB
```

The resulting output:

No feedback

An invocation example

```
chrcrelationship -cycleperiodseconds 20 relC
```


The resulting output:

No feedback

lspartnership

Use the **lspartnership** command to display a concise or detailed view of the current clustered systems (systems) that are associated with the local system.

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 displayed. If a capacity is specified, the units must also be included.

- 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 command-line interface (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 (""): `lspartnership -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.

-filtervalue?

(Optional) displays a list of filters that can be applied against this view. The following filter attributes are valid:

- id
- name

system_id | system_name

(Optional) Specifies the name or ID of a system. Using this parameter displays the detailed view of the specific partner system, and any value specified by the **-filtervalue** (which filters a view that is based on specific attribute values that relate to each object type) parameter is ignored. When specifying *system_id* or *system_name* parameter, the concise view of all systems that match the filtering requirements that are specified by the **-filtervalue** parameter are displayed.

Description

Table 92 described attribute values.

Table 92. lspartnership attribute values

Attribute	Value
id	Indicates the system ID.
name	Indicates the system name.
location	Indicates the system location.
code_level	Indicates the code level.

Table 92. *lspartnership* attribute values (continued)

Attribute	Value
partnership	<p>Indicates the current state of the partnership; not applicable for the local system and is blank.</p> <p>The partnership field can show the following values:</p> <p>fully_configured The mkfcpartnership or mkippartnership command has been issued in both directions and the remote system is online and available.</p> <p>partially_configured_local The mkfcpartnership or mkippartnership command has only been issued from the local system to the remote system. The remote system is online and available for partnership.</p> <p>partially_configured_local_stopped The mkfcpartnership or mkippartnership command has only been issued from the local system to the remote system. The chpartnership command with the stop parameter has been issued from the local system, and the remote system is online and available. Issue chpartnership -start on the local system, and mkfcpartnership or mkippartnership on the remote system</p> <p>not_present The mkfcpartnership or mkippartnership command has been issued from the local system to the remote system, and the remote system is not available. The remote system is either offline or not connected to the local system.</p> <p>fully_configured_stopped The mkfcpartnership or mkippartnership command has been issued in both directions and the remote system is online and available. The chpartnership command with the stop parameter has been issued from the local system.</p> <p>fully_configured_remote_stopped The mkfcpartnership or mkippartnership command has been issued in both directions and the remote system is online and available. The chpartnership command with the stop parameter has been issued from the remote system.</p> <p>fully_configured_local_excluded The mkfcpartnership or mkippartnership command has been issued in both directions. The local system excludes the connection to the remote system and the partnership is unable to sustain the I/O workload for the Metro Mirror or Global Mirror relationships.</p> <p>fully_configured_remote_excluded The mkfcpartnership or mkippartnership command has been issued in both directions. The local system excludes the connection to the remote system and the partnership is unable to sustain the I/O workload for the Metro Mirror or Global Mirror relationships.</p> <p>fully_configured_exceeded There are too many systems in the system network and the partnership from the local system to the remote has been disabled.</p>
relationship_bandwidth_limit	Indicates the current bandwidth limit.
type	<p>Indicates the type of partnership:</p> <ul style="list-style-type: none"> • Fibre Channel (FC) • Internet Protocol Version 4 (IPv4) or Internet Protocol Version 6 (IPv6) <p>FC partnerships are created connecting two systems over FC or Fibre Channel over Ethernet (FCoE) fabrics. IPv4 or IPv6 partnerships are created connecting two systems over native IP links.</p>

Table 92. *lspartnership* attribute values (continued)

Attribute	Value
cluster_ip	Indicates the partner system IP address, which can be IPv4 or IPv6. This information is displayed for IP-based partnerships. For IP-based partnership this field displays the system IP address specified while creating the partnership using mkippartnership .
chap_secret	Indicates the Challenge-Handshake Access Protocol (CHAP) secret (up to eighty alphanumeric characters) for the partner system. The CHAP authenticates the local system with the partner system during discover and Internet Small Computer System Interface (iSCSI) system session creation. For FC-based and FCoE-based relationships this field is always blank.
link_bandwidth_mbits	Indicates the aggregate bandwidth for the Remote Copy (RC) link in megabits per second (Mbps). This is a numeric value from 0 to 100000. If there are multiple links between the local and remote systems, this parameter is set to the sum of the link bandwidths for these links.
background_copy_rate	Indicates the bandwidth allocation for background copy operations performed over the replication link. It is expressed as a percentage of the link bandwidth value, and is the maximum rate at which background copy operations are performed. This is a numeric value from 0 to 100 .
event_log_sequence	Indicates the last sequence number (indicating the last event) from event log for this partnership. This is a numeric value from 100 to 8000000. For FC-based and FCoE-based relationships this field is always blank.
max_replication_delay	Indicates the value for maximum replication delay. This is a numeric value from 0 to 360.
compressed	Indicates whether compression is enabled. The values are yes and no (default).

A concise invocation example

```
lspartnership
```

The concise resulting output:

```
id          name      location  partnership          type  cluster_ip          event_log_sequence
000002006BC0A0D4  system-1  local    partially_configured_local  ipv6  fe80::200:f8ff:fe21:67cf
000002006200A0E5  system-2  remote   partially_configured_local  fc
000002006200A0F6  system-3  remote   partially_configured_local  fc
000002006200A0G7  system-4  remote   partially_configured_local  fc
```

A detailed invocation example

```
lspartnership cluster-2
```

The detailed resulting output:

```
id 000002006200A0EA
name system-2
location remote
partnership partially_configured_local
code_level 6.3.0.0 (build 35.7.1105071000)
console_IP 9.180.28.63:443
gm_link_tolerance 300
gm_inter_system_delay_simulation 0
gm_intra_system_delay_simulation 0
relationship_bandwidth_limit 25
gm_max_host_delay 5
type fc
cluster_ip
chap_secret
event_log_sequence
```

```
link_bandwidth_mbits 1024
background_copy_rate 25
max_replication_delay 145
compressed yes
```

lspartnershipcandidate

Use the **lspartnershipcandidate** command to list the clustered systems available for setting up a partnership with the local system. This is a prerequisite for creating inter-system Metro or Global Mirror relationships.

Syntax

```
►► — lspartnershipcandidate — [ -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 systems that are available as candidate partner systems to form a Metro Mirror or Global Mirror partnership between two systems.

Output from the command shows the system ID, name, and configured status of the remote candidate system. The remote candidate system forms a partnership with the local system when you use the **mkippartnership** or **mkfcpartnership** command. The remote system shows the partnership status as **partially_configured_local_stopped** or **partially_configured_local** when you use the **lssystem** command. The **lspartnershipcandidate** command displays the configured status of those remote systems that have formed a partnership with the local system.

An invocation example

```
lspartnershipcandidate
```


The resulting output:

id	configured	system_name
0000010034E0F430	no	ldsystem26

lsrconsistgrp

Use the **lsrconsistgrp** command to return a concise list or a detailed view of remote copy relationships such as Metro Mirror, Global Mirror, or active-active consistency groups visible to the system.

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 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 command line interface (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:

```
lsrconsistgrp -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 **lsrconsistgrp** 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 remote copy relationships such as Global Mirror, Metro Mirror, or active-active consistency groups that are visible to the system.

Table 93 provides possible values for the attributes that are displayed as data in the output views.

Table 93. lsrrconsistgrp command output values

Attribute	Value
primary	n/a, master, aux
state	The state: <ul style="list-style-type: none"> • consistent_copying • inconsistent_stopped • inconsistent_copying • consistent_stopped • consistent_synchronized • idling • idling_disconnected • inconsistent_disconnected • consistent_disconnected • empty
cycle_period_seconds	The minimum period in seconds between multiple cycles (integer between 60 and 86400; default is 300).
cycling_mode	The type of Global Mirroring, Metro Mirroring, or active-active cycling to use: none (default) or multi
freeze_time	The time in YYMMDDHHMM format

Table 93. *lsrconsistgrp* command output values (continued)

Attribute	Value
status	<p>The values are:</p> <ul style="list-style-type: none"> • online, which indicates that the relationship is online and accessible. If the relationship state is ConsistentSynchronized, ConsistentCopying, or InconsistentCopying the volumes replicate host I/O operations that the primary volume receives. • primary_offline, which indicates that the primary volume in the relationship is offline. This prevents additional I/O operations, and synchronization pauses until the primary volume is online again. • secondary_offline, which indicates that the secondary volume in the relationship is offline. For r Global Mirror relationships in a ConsistentSynchronized state (no change volumes) and Metro Mirror relationships, additional I/O write operations to the primary volume terminate the relationship. • io_channel_offline, which indicates that the remote system is not accessible. For Global Mirror relationships in a ConsistentSynchronized state (no change volumes) and Metro Mirror relationships, additional I/O write operations to the primary volume terminate the relationship. • primary_change_offline, which indicates that the primary change volume in the relationship is offline. For Global Mirror with change volume relationships, the current I/O cycle ends, and a new I/O cycle begins when the primary change volume is online again. • secondary_change_offline, which indicates that the secondary change volume in the relationship is offline. For Global Mirror with change volume relationships, the current I/O cycle pauses and a new I/O cycle resumes when the secondary volume is online again. • change_volumes_needed, indicates an active-active relationship that is in a HyperSwap volume or Global Mirror volume with a change volume relationship. Additionally, at least one change volume is not configured. <p>Important: Replication services are not usable.</p> <p>Remember: This field is blank.</p>
sync	<p>The values for this field:</p> <ul style="list-style-type: none"> • in_sync • out_of_sync
mutually_consistent	<p>The values for this field:</p> <ul style="list-style-type: none"> • yes • no • blank <p>Note: This relationship is consistent with other consistency group relationships. This value is blank unless there is also a value in consistent_stopped, consistent_disconnected, and consistent_copying.</p>
copy_type	<p>The copy type:</p> <ul style="list-style-type: none"> • metro • global • activeactive • blank

Note: The names of the Global Mirror or Metro Mirror relationships and consistency groups might be blank if the relationship or consistency groups are intersystem and the system partnership is disconnected.

The sync attribute has a value of `in_sync` when the contents are synchronized (identical) between 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

```
lsrcconsistgrp -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:cycling_mode:freeze_time

 248:jdemo_BA_cons1:0000020060406746:clusterB:0000020061413ABA:clusterA:master:
consistent_stopped:2:global:none:06/06/27/08/31/37
 249:rccstgrp0:0000020061413ABA:clusterA:0000020061413ABA:clusterA::empty:0
:empty_group
 250:jdemo_BA_cons2:0000020060406746:clusterB:0000020061413ABA:clusterA:master:
inconsistent_stopped:1:metro:none:06/06/27/08/31/37
 251:BA_cons1:0000020060406746:clusterB:0000020061413ABA:clusterA:master:
consistent_stopped:4:metro:none:06/06/27/08/31/37
 252:AB_cons2:0000020061413ABA:clusterA:0000020060406746:clusterB::empty:0
:empty_group:none:06/06/27/08/31/37
 253:AB_cons1:0000020061413ABA:clusterA:0000020060406746:clusterB:aux:
consistent_stopped:3:global:none:06/06/27/08/31/37
 254:AA_cons2:0000020061413ABA:clusterA:0000020061413ABA:clusterA::empty:0
:empty_group:none:06/06/27/08/31/37
 255:AA_cons1:0000020061413ABA:clusterA:0000020061413ABA:clusterA:master:
consistent_synchronized:2:global:none:06/06/27/08/31/37
```

A detailed invocation example

```
lsrcconsistgrp -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:consistent_synchronized
relationship_count:1
freeze_time:06/06/27/08/31/37
status:online
sync:in_sync
copy_type:activeactive
cycle_period_seconds:300
cycling_mode:none
RC_rel_id:2
RC_rel_name:aaa
```

lsrcrelationship

Use the **lsrcrelationship** command to return a concise list or a detailed view of remote copy relationships such as Metro Mirror, Global Mirror or active-active relationships visible to the system.

Syntax

```
➤— lsrcrelationship — [ -filtervalue — attribute=value ] [ -nohdr ] —————➤
```



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:

```
lsrcrelationship -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 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 1-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 that match 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 **lsrcrelationship** command are:

- RC_rel_id
- RC_rel_name
- master_system_id
- master_system_name
- master_vdisk_id
- master_vdisk_name
- aus_system_id
- aus_system_name
- aus_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 remote copy relationships such as Metro Mirror, Global Mirror, or active-active relationships visible to the system.

Table 94 provides possible values for the attributes that are displayed as data in the output views.

Table 94. lscrelationship command attributes and values

Attribute	Value
primary	n/a, master, aux
state	consistent_copying, inconsistent_stopped, inconsistent_copying, consistent_stopped, consistent_synchronized, idling, idling_disconnected, inconsistent_disconnected, consistent_disconnected
progress	Integer between 0 and 100, n/a
cycle_period_seconds	The minimum period in seconds between multiple cycles (integer between 60 and 86400; default is 300).
cycling_mode	The type of Global Mirror, Metro Mirroring, or active-active cycling to use. The choices are: <ul style="list-style-type: none"> • none (default) • multi
copy_type	The copy type: <ul style="list-style-type: none"> • metro • global • activeactive • blank
freeze time	The time in YY/MM/DD/HH/MM format

Table 94. `lsrcrelationship` command attributes and values (continued)

Attribute	Value
status	<p>The values are:</p> <ul style="list-style-type: none"> • <code>online</code>, which indicates that the relationship is online and accessible. If the relationship state is <code>ConsistentSynchronized</code>, <code>ConsistentCopying</code>, or <code>InconsistentCopying</code> the volumes replicate host I/O operations that the primary volume receives. • <code>primary_offline</code>, which indicates that the primary volume in the relationship is offline. This prevents additional I/O operations, and synchronization pauses until the primary volume is online again. • <code>secondary_offline</code>, which indicates that the secondary volume in the relationship is offline. For <code>r</code> Global Mirror relationships in a <code>ConsistentSynchronized</code> state (no change volumes) and Metro Mirror relationships, additional I/O write operations to the primary volume terminate the relationship. • <code>io_channel_offline</code>, which indicates that the remote system is not accessible. For Global Mirror relationships in a <code>ConsistentSynchronized</code> state (no change volumes) and Metro Mirror relationships, additional I/O write operations to the primary volume terminate the relationship. • <code>primary_change_offline</code>, which indicates that the primary change volume in the relationship is offline. For Global Mirror with change volume relationships, the current I/O cycle ends, and a new I/O cycle begins when the primary change volume is online again. • <code>secondary_change_offline</code>, which indicates that the secondary change volume in the relationship is offline. For Global Mirror with change volume relationships, the current I/O cycle pauses and a new I/O cycle resumes when the secondary volume is online again. • <code>change_volumes_needed</code>, indicates an active-active relationship that is in a HyperSwap volumes or Global Mirror with change volume relationships. Additionally, at least one change volume is not configured. <p>Important: Replication services are not usable.</p>
sync	n/a, <code>in_sync</code> , <code>out_of_sync</code>
master_change_vdisk_name	<p>The name of the volume acting as the master change volume for the relationship (blank if not defined)</p> <p>Note: This field identifies the change volume for the master volume if configured. For an intersystem relationship, if the master volume is in the other clustered system (system), the master change volume is also in the other system.</p>
aux_change_vdisk_id	<p>The id of the volume acting as the auxiliary change volume for the relationship (blank if not defined)</p> <p>Note: This field identifies the change volume for the auxiliary volume, if such a volume has been configured. For an intersystem relationship, if the auxiliary volume is in the other system, the auxiliary change volume is also in the other system.</p>
aux_change_vdisk_name	<p>The name of the volume acting as the auxiliary change volume for the relationship (blank if not defined)</p> <p>Note: This field identifies the change volume for the auxiliary volume if configured. For an intersystem relationship, if the auxiliary volume is in the other system, the auxiliary change volume is also in the other system.</p>
bg_copy_priority	Unused.

Note: The names of the Global Mirror, Metro Mirror, or active-active relationships and consistency groups can be blank if the relationship or consistency groups are intersystem and the system partnership is disconnected.

The `change_volumes_needed` status is set if either the master or auxiliary change volume is not defined for relationships with either of these:

- `copy_type` set to `activeactive`
- `copy_type` set to `global` and `cycling_mode` set to `multi`

The `sync` attribute has a value of `in_sync` when the contents are synchronized (identical) between 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

```
lsrcrelationship -delim : -filtervalue name=j*
```

The concise resulting output:

```
id: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:bg_copy
_priority:progress:copy_type:cycling_mode:freeze_time
45:jrel_AB1:0000020061413ABA:clusterA:45:jdisk_B8:0000020060406746:clusterB:38:j
disk_B1:master:::consistent_stopped:50:metro:none:06/06/27/08/31/37
48:jrel_AB2:0000020061413ABA:clusterA:48:jdisk_A4:0000020060406746:clusterB:41:j
disk_B4:master:::consistent_synchronized:50:metro:none:06/06/27/09/31/37
49:jrel_AB_1:0000020060406746:clusterB:42:jdisk_B5:0000020061413ABA:clusterA:49:j
disk_A5:master:248:jdemo_BA_cons1:consistent_stopped:50:metro:none:06/06/27/10/31/37
50:jrel_AB_2:0000020060406746:clusterB:43:jdisk_B6:0000020061413ABA:clusterA:
50:jdisk_A6:master:248:jdemo_BA_cons1:consistent_stopped:50:metro:none:06/06/27/11/31/37
```

A detailed invocation example

```
lsrcrelationship -delim : AB_2
```


The detailed resulting output:

```
id:9
name:AB_2
master_cluster_id:0000020061413MOE
master_cluster_name:chelseaB
master_vdisk_id:9
master_vdisk_name:stripe9
aux_cluster_id:0000020061413MOE
aux_cluster_name:chelseaB
aux_vdisk_id:10
aux_vdisk_name:stripe9_b
copy_type:activeactive
cycle_period_seconds:300
cycling_mode:multi
primary:master
consistency_group_id:
consistency_group_name:
state:consistent_synchronized
bg_copy_priority:50
progress:
freeze_time:2006/05/05/08/26/46
status:online
sync:in_sync
```

lsrcrelationshipcandidate

Use the `lsrcrelationshipcandidate` command to lists volumes that can form Metro Mirror, Global Mirror or active-active relationships. You can list eligible volumes that are on the local or remote clustered system (system).

Syntax



Parameters

-master *master_vdisk_id* | *master_vdisk_name*

(Required) Specifies a particular volume to use as the master volume. The command finds candidates that match the size of this volume. If you are requesting candidate volumes on the local system, this command also matches the io_group.

-aux *aux_cluster_id* | *aux_cluster_name*

(Required) Specifies a remote system with volume candidates for an intersystem relationship. If you do not specify this parameter, the candidates on the local system 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 volumes that can be either the master or the auxiliary disk for a Metro Mirror, Global Mirror, or active-active relationship. Volume IDs and names are displayed.

Note: Volumes that are flash disks are excluded from the view when a FlashCopy map is constructed.

An invocation example

```
lsrelationshipcandidate -delim :
```

The resulting output:

```
id:vdisk_name
0:vdisk0
4:vdisk4
```

lsrelationshipprogress

Use the **lsrelationshipprogress** command to display the progress of the background copy of a Metro Mirror, Global Mirror, or active-active relationship as a percentage. When the initial background copy process for a relationship completes, a null value is displayed for the progress of that relationship.

Syntax

```
lsrcrelationshipprogress - [-nohdr] [-delim delimiter]
[rcrelationship_id
rcrelationship_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.

rcrelationship_id | rcrelationship_name

(Required) Specifies the object ID or name of the specified type.

Description

This command displays the progress of the background copy of a Metro Mirror, Global Mirror, or active-active relationship as a percentage.

An invocation example

```
lsrcrelationshipprogress -delim : 0
```

The resulting output:

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

mkfcpartnership

Use the **mkfcpartnership** command to define partnerships using Fibre Channel (FC) or Fibre Channel over Ethernet (FCoE).

Syntax

```
mkfcpartnership --linkbandwidthmbits link_bandwidth_in_mbps
[-backgroundcopyrate percentage] [remote_system_id
remote_system_name]
```


Parameters

-linkbandwidthmbits *link_bandwidth_in_mbps*

(Required) Specifies the aggregate bandwidth of the Remote Copy (RC) link between two clustered systems (systems in megabits per second (Mbps)). It is a numeric value from 1 to 100000.

Note: Remote copy includes Metro Mirror and Global Mirror.
This value remains the same after a system update.

Note: If the specified value is non-zero, the combination of both the **-backgroundcopyrate** and the **-linkbandwidthmbits** values must result in a background copy bandwidth of at least 8 Mbps.

-backgroundcopyrate *percentage*

(Optional) Specifies the maximum percentage of aggregate link bandwidth that can be used for background copy operations. It is a numeric value from 0 to 100, and the default value is 50, which means that a maximum of 50% of the aggregate link bandwidth can be used for background copy operations. This parameter can be specified without stopping the partnership.

Note: If the specified value is non-zero, the combination of both the **-backgroundcopyrate** and the **-linkbandwidthmbits** values must result in a background copy bandwidth of at least 8 Mbps.

remote_system_id | *remote_system_name*

(Required) Specifies the remote system ID or name of a partnership. The specified value must match one of the system IDs or names returned after issuing **lspartnershipcandidate**.

Description

This command defines FC-based or FCoE-based partnerships. However, all existing partnerships are automatically updated to FC partnerships, any invocation of this command is applicable only to FC-based partnerships, and all partnerships created are FC-based partnerships.

An invocation example

```
mkfcpartnership -linkbandwidthmbits 100 -backgroundcopyrate 50 remote-system-2
```

The resulting output:

No feedback

An invocation example

```
mkfcpartnership -linkbandwidthmbits 1024 -backgroundcopyrate 25 remote-system-3
```

The resulting output:


No feedback

mkippartnership

Use the **mkippartnership** command to define a new partnership created over Internet Protocol (IP) links.

Syntax

```
►►—mkippartnership— --type— ipv4 ipv6 -clusterip—ipadr— -chapsecret—CHAPsecret— →  
►—linkbandwidthmbits—link_bandwidth_in_mbps— -backgroundcopyrate—percentage— →
```



Parameters

-type *ipv4* | *ipv6*

(Required) Specifies the Internet Protocol (IP) address format for the partnership using either of the following case-sensitive strings:

- *ipv4* for Internet Protocol Version 4 (IPv4)
- *ipv6* for Internet Protocol Version 6 (IPv6)

All Transmission Control Protocol (TCP) Remote Copy (RC) connections between the primary and remote clustered systems (systems) are created using specific IP addresses.

Note: Remote copy includes Metro Mirror and Global Mirror.

Partnership creation fails if the Internet Protocol (IP) address types specified for either primary or remote systems are not the same.

-clusterip *ipadr*

(Required) Specifies the partner system IP address, either *ipv4* or *ipv6*. Systems connected over IP links are not displayed by **lspartnershipcandidate** before executing **mkippartnership**. This does not apply to FC-based or FCoE-based connections.

-chapsecret *CHAPsecret*

(Optional) Specifies the Challenge-Handshake Authentication Protocol (CHAP) secret of the partner system. The maximum size of the CHAP secret is eighty alphanumeric characters.

-linkbandwidthmbits *link_bandwidth_in_mbps*

(Required) Specifies the aggregate bandwidth of the RC link between two clustered systems (systems) in megabits per second (Mbps). It is a numeric value from 1 to 100000. This parameter can be specified without stopping the partnership.

Note: If the specified value is non-zero, the combination of both the **-backgroundcopyrate** and the **-linkbandwidthmbits** values must result in a background copy bandwidth of at least 8 Mbps.

-backgroundcopyrate *percentage*

(Optional) Specifies the maximum percentage of aggregate link bandwidth that can be used for background copy operations. It is a numeric value from 0 to 100, and the default value is 50, which means that a maximum of 50% of the aggregate link bandwidth can be used for background copy operations.

Note: If the specified value is non-zero, the combination of both the **-backgroundcopyrate** and the **-linkbandwidthmbits** values must result in a background copy bandwidth of at least 8 Mbps.

-compressed *yes* | *no*

(Optional) Specifies whether compression is enabled for this partnership. The default value is *no*.

Description

This command defines a new partnership created over Internet Protocol (IP) links. A remote system IP must be specified so its IP ports are enabled for data replication. RC sessions can then be created between the two partners.

In FC-based or FCoE-based partnerships, the partner system must first be a partnership candidate (displayed by **lspartnership**). Then it can become part of a partnership, created specifying **mkfcpartnership** with the remote system ID or name.

The **lspartnershipcandidate** command displays partnership candidates.

For IP partnerships, specifying **mkippartnership** with the cluster IP address and CHAP secret of the partner creates the partnership.

All TCP connections are established using either IPv4 or IPv6, and it cannot be a mix of the two IP address types.

Both systems in a partnership must have at least one IP address from an identical replication group to establish RC partnerships. Replication groups are numeric values that specify the pools of local IP addresses that establish Remote Copy partnerships with pools of IP addresses configured on the partner system.

An invocation example

```
mkippartnership -type ipv4 -clusterip 192.168.32.19
                -chapsecret mychapsecret -linkbandwidthmbits 100 -backgroundcopyrate 50
```

The resulting output:

No feedback

An invocation example

```
mkippartnership -type ipv6 -clusterip fe80::200:f8ff:fe21:67cf
                -chapsecret mychapsecret -linkbandwidthmbits 1024 -backgroundcopyrate 25
```

The resulting output:

No feedback

An invocation example creating a compressed IP replication link

```
mkippartnership -type ipv4 -clusterip 192.168.32.19 -chapsecret mychapsecret -linkbandwidthmbits 100 -backgroundcopyrate 50 -comp
```

The resulting output:

No feedback

mkpartnership (Discontinued)

The **mkpartnership** command is deprecated. Use either the **mkfcpartnership** or **mkippartnership** command instead.

mkrcconsistgrp

Use the **mkrcconsistgrp** command to create a new, empty remote copy consistency group. If the **-cluster** parameter is not specified, the consistency group is created on the local clustered system (system) only.

Syntax

```
➤➤ 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 system. If **-cluster** is not specified, a consistency group is created only on the local system.

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 systems within this consistency group. If the consistency group involves two system, the systems must be in communication throughout the create process.

The new consistency group does not contain any relationships and is in an empty state. You can add Metro Mirror, Global Mirror, or active-active relationships to the group using the **chrcrelationship** command.

Remember: Names representing remote copy consistency groups relationships are restricted to fifteen characters in length (not sixty-three for an extended character set).

An invocation example

```
mkrconsistgrp -name rc_testgrp
```

The resulting output:

```
RC Consistency Group, id [255], successfully created
```

mkrcrelationship

Use the **mkrcrelationship** command to create a new Global Mirror, Metro Mirror, or active-active relationship with volumes in the same clustered system (system), forming an intrasystem Metro Mirror relationship or intersystem relationship (if it involves more than one system).

Syntax

```
►► mkrcrelationship — — -master ———— master_vdisk_id ———— -aux ———— aux_vdisk_id ————
                               └── master_vdisk_name ───┘           └── aux_vdisk_name ───┘
► -cluster ———— cluster_id ———— ┬── -name ———— new_name_id ───┘
                               └── cluster_name ───┘
► ┬── -consistgrp ———— consist_group_id ———— ┬── -sync ───┘
  │   └── consist_group_name ───┘
► ┬── -global ———— ┬── -cyclingmode ———— none ───┘
  │   └── ┬── ───┘ └── multi ───┘
  └── -activeactive ───┘
```

Parameters

-master *master_vdisk_id* | *master_vdisk_name*

(Required) Specifies the ID or name of the *master_vdisk_id* or *master_vdisk_name*.

If a new remote copy relationship is currently mapped to a host of type *hide_secondary*, the secondary volume is not presented to the host; however, it is mapped for configuration purposes. The secondary volume is presented to the host if the:

- Host type is changed to a type other than `hide_secondary`
- Remote copy relationship is stopped by specifying **-access**
- Volume is no longer a secondary volume because the remote copy relationship is deleted or switched

-aux *aux_vdisk_id* | *aux_vdisk_name*

(Required) Specifies the ID or name of the `aux_vdisk_id` or `aux_vdisk_name`.

-cluster *cluster_id* | *cluster_name*

(Required) Specifies the ID or name of the remote cluster.

- If you are creating an intrasystem relationship, enter the ID of the local system. The volumes in the relationship must belong to the same I/O group within the system.
- If you are creating an intersystem relationship, enter the ID of the remote system. To create a relationship in two different systems, the systems must be connected at the time that the **mkrcrelationship** 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 Mirror, Global Mirror, and active-active 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. You cannot specify this keyword with **-activeactive**.

-cyclingmode *none* | *multi*

(Optional) Specifies the behavior of Global Mirror for this relationship.

- Specifying *none*, the default, gives identical behavior to Global Mirror in previous versions of SAN Volume Controller.
- Specifying *multi* uses the cycling protocol.

The default cycle period is 300 seconds. The cycle period can be modified after the relationship is created by using the **chrcrelationship** command. To start a relationship with `cycling_mode` set to *multi*, change volumes must be defined for the relationship.

Important: This parameter must be specified with **-global**.

-activeactive

(Optional) Specifies that the relationship is created in an active-active mode. You cannot specify this keyword with **-global** (this parameter defaults to a Metro Mirror relationship being created).

Description

This command creates a new Global Mirror, Metro Mirror, or active-active relationship. A Metro Mirror relationship defines the relationship between two volumes. One volume is a master volume and the other volume is an auxiliary volume. This relationship persists until deleted. The auxiliary volume must be identical in size to the master volume or the command fails. This command also returns the new relationship ID.

The master and auxiliary cannot be in an existing relationship. Any defined FlashCopy mappings that have the proposed master volume as the target of the FlashCopy mapping must be using the same I/O group as the master volume. Any defined FlashCopy mappings that have the proposed auxiliary volume as the target of the FlashCopy mapping must be using the same I/O group as the auxiliary volume.

Note: You cannot create a remote copy relationship with this command if the auxiliary volume is an active FlashCopy mapping target. If the I/O group has enough bitmap space available to allocate for remote copy and the allocated space for the remote copy is not large enough to accommodate the new relationship, space is automatically added. (Remote copy includes Global Mirror, Metro Mirror, and active-active relationships.)

Note: You cannot use this command if cloud snapshot is enabled on the volume or the volume owner type is `cloud_backup`.

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 volumes before sending confirmation of I/O completion to the host application. This ensures that the secondary volume is synchronized with the primary volume if 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 volume. If a failover operation is performed, the host application must recover and apply any updates that were not committed to the secondary volume.

You can optionally give the relationship a name. The name must be a unique relationship name across both systems.

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 volumes 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 volumes will stop in a consistent manner if the relationships that are associated with the volumes 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 volume and the data volume at the secondary site are in a consistent state. This can be achieved by putting the relationships that are associated with these volumes into a consistency group. Both Metro Mirror and Global Mirror processing ensure that updates to both volumes at the secondary 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 system and the same auxiliary system. 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 volumes 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 volume before you issue the **mkrcrelationship** command.

If you specify the **-global** parameter, a Global Mirror relationship is created. Otherwise, a Metro Mirror relationship is created instead.

The volumes that are specified on the **-master** and **-aux** parameters cannot be master or auxiliary volumes in an existing relationship.

If you specify **-activeactive**:

- The system specified with **-cluster** must be the local system.
- **-global** must not be specified.
- The volume specified with **-master** must:
 - Be in an I/O group with both nodes having the same site name and site ID
 - Have all volume copies stored in storage pools in the same site as the volume's I/O group
 - Not be the target of a FlashCopy mapping
 - Not be the source of any FlashCopy mappings to volumes in a different site or using bitmap memory from nodes in a different site (but the volume can be the source of a FlashCopy mapping in which the target volume and map are in the same site)
- The volume that is specified with **-aux** must:
 - Be in an I/O group with a different site ID and site name to the master volume, and must not have any volume host mappings defined
 - Have all volume copies stored in storage pools in the same site as the volume's I/O group
 - Not be the target of a FlashCopy mapping
 - Not be the source of any FlashCopy mappings to volumes in a different site or using bitmap memory from nodes in a different site (but the volume can be the source of a FlashCopy mapping in which the target volume and map are in the same site)

Access the data stored on these volumes by accessing the volume specified by the **-master** parameter. Both I/O groups of the volumes specified by the **-master** and **-aux** parameters have a local physical copy and cache, allowing access (using the master volume ID) whether the auxiliary volume's site is available or not.

Remember: This command cannot be used on a volume that is owned by a file system.

An invocation example

```
mkrcrelationship -master vdisk1 -aux vdisk2 -name rccopy1  
-cluster 0000020063432AFD
```

The resulting output:

```
RC Relationship, id [28], successfully created
```

An invocation example

```
mkrcrelationship -master vdiskA -aux vdiskB -cluster clusterB -name new_rel -global -cyclingmode multi
```

The resulting output:

```
RC Relationship, id [28], successfully created
```

An invocation example

```
mkrcrelationship -master volA -aux volB -cluster localCluster -activeactive
```

The resulting output:

```
RC Relationship, id [28], successfully created
```

rmpartnership

Use the **rmpartnership** command to remove a Metro Mirror or Global Mirror partnership on one clustered system (system). Because the partnership exists on both systems, it is necessary to run this command on both systems to remove both sides of the partnership. If the command is run on only one system, the partnership enters a partially configured state on the other system.

Syntax

```
➤— rmpartnership —┬— remote_cluster_id —┐—————➤
                   └— remote_cluster_name —┘
```

Parameters

remote_cluster_id | *remote_cluster_name*
(Required) Specifies the system ID or the name of the remote system.

Description

This command deletes one half of a partnership on a system. To remove the entire partnership, you must run the command twice, once on each system.

Attention: Before running the **rmpartnership** command, you must remove all relationships and groups that are defined between the two systems. To display system relationships and groups, run the **lsrcrelationship** and **lsrconsistgrp** commands. To remove the relationships and groups that are defined between the two systems, run the **rmrcrelationship** and **rmrconsistgrp** commands.

An invocation example

```
rmpartnership cluster1
```

The resulting output:

```
No feedback
```

rmrconsistgrp

Use the **rmrconsistgrp** command to delete an existing Metro Mirror, Global Mirror, or active-active consistency group.

Syntax

```
➤➤ rmrconsistgrp — [ -force ] [ rc_consist_group_id | rc_consist_group_name ] ➤➤
```

Parameters

-force

(Optional) Specifies that you want the system to remove any relationship belonging to a group before the consistency group is deleted. The relationship itself is not deleted; it becomes a stand-alone relationship.

Note: The **-force** parameter must be used to delete a consistency group when the consistency group has any Metro Mirror, Global Mirror, or active-active relationship 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 **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

```
rmrconsistgrp rctestone
```

The resulting output:

No feedback

rmrrelationship

Use the **rmrrelationship** command to delete an existing remote copy relationship.

Syntax

```
➤➤ rmrrelationship — [ -force ] [ rc_rel_id | rc_rel_name ] ➤➤
```

Parameters

-force

(Optional) Specifies that the relationship should be deleted even if it results in the secondary volume containing inconsistent data. This only applies to active-active relationships or Global Mirror relationships using multi cycling mode.

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 volumes. It does not affect the volumes themselves.

If the relationship is disconnected at the time that the command is issued, the relationship is only deleted on the clustered system (system) where the command is being run. When the systems reconnect, the relationship is automatically deleted on the other system. Alternatively, if the systems are disconnected and if you still want to remove the relationship on both systems, you can issue the **rmrcrelationship** command independently on both of the systems.

If a relationship is active-active or a Global Mirror relationship using multicycling mode, and you attempt to delete the relationship without enabling access first, specifying **rmrcrelationship** might fail with an error because the relationship does not currently have a fully consistent secondary volume. Specifying **-force** overrides this test. This is not the default behavior, and you can quiesce and delete the relationship in order to use the secondary volume's data immediately. If the map is still performing the background copy to migrate data from the change volume to the secondary volume, the changed volume and associated FlashCopy mappings remain defined when **rmrcrelationship** completes. The FlashCopy mappings are deleted after the background copy completes, and the change volume becomes unusable again.

If you delete an inconsistent relationship, the secondary volume becomes accessible even though it is still inconsistent. This is the one case in which Metro Mirror, Global Mirror or HyperSwap does not inhibit access to inconsistent data.

An invocation example

rmrcrelationship rccopy1


The resulting output:

No feedback

startrcconsistgrp

Use the **starttrconsistgrp** command to start the Global Mirror, Metro Mirror, or active-active consistency group copy process, set the direction of copy if it is undefined, and optionally mark the secondary volumes of the consistency group as clean.

Syntax



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 - except if consistency protection is configured.

-clean

(Optional) Specifies that the volume that is to become a secondary is clean for each of the relationships belonging to the group; any changes made on the secondary volume are ignored, and only changes made on the clean primary volume 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 if all data changed on the secondary volumes while the consistency group was in the `idling` state matches the state of the primary volumes when the consistency group was stopped. Otherwise, relationships that are not consistent are reported as consistent. Once this has been done there is no method to determine whether these volumes 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 Global Mirror, Metro Mirror, or active-active stand-alone consistency group. You cannot use this command to start a remote copy relationship if the primary volume is a target volume of a prepared FlashCopy mapping.

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.

Note: You cannot start a consistency group if it contains a relationship with primary and secondary volumes that are different sizes.

If an existing remote copy relationship is stopped with **-access** specified but is restarted and the resultant secondary volume (depending on the choice of primary) is mapped to a host of type `hide_secondary`, that volume is not presented to the host. This is true even though it is mapped for configuration purposes. The mapped volumes are presented to the host if the:

- Host type is changed to a type other than `hide_secondary`
- Remote copy relationship is stopped and **-access** is specified
- Volume ceases to be a secondary volume because the remote copy relationship is being deleted or switched

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 volumes 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

The command does not fail if you specify the **-force** parameter.

When you configure a secondary change volume on all relationships in a consistency group, the consistency group goes into a `consistent_copying` state during resynchronization. If you specify **starttrcconsistgrp** for a consistency group in an idling state, consistency protection is disabled if any of the secondary volumes have been written to. This means that you must specify the **-force** parameter.

The **-clean** parameter is used when a Global Mirror or Metro Mirror group is started and the secondary volumes in this group are assumed to be clean, which 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 volumes. 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 **stoptrcconsistgrp** 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 **starttrcconsistgrp** 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.

After restarting a consistency group in either of these states (Idling or multi), the data on the secondary volumes is not usable for disaster recovery until the consistency group becomes consistent.

A Global Mirror consistency group with a `cycling_mode` of `multi` in either of these states does not require the **-force** parameter because consistent secondary images are retained. However, if such a consistency group is in idling state and written data has been received at any secondary volume in the consistency group, the **-force** flag is still required, because the secondary volumes have a divergent image that cannot represent a consistent earlier state.

A Global Mirror relationship with a cycling mode of:

- *none* uses the non-cycling Global Mirror algorithm
- *multi* must have a change volume configured at the primary volume (or the command fails)
- *multi* must also have a change volume configured at the secondary volume (or the command fails)
- *multi* performs multiple cycles of cycling

After creating a background copy the relationship remains in copying state, waiting for the remainder of the period time to expire before performing a new cycle. If the secondary change volume is deconfigured when the background copy completes, the relationship stops as if there is no cycle period.

Relationships that are active-active must have a state of idling to be started.

An invocation example

```
starttrconsistgrp rccopy1
```


The resulting output:

No feedback

starttrrelationship

Use the **starttrrelationship** command to start the Metro Mirror or Global Mirror relationship copy process, set the direction of copy if undefined, and (optionally) mark the secondary volume of the relationship as clean. The relationship must be a stand-alone relationship. You can also use this command to restart the active-active relationship copy process after specifying **stoptrrelationship -access**.

Syntax



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 idling state - except if consistency protection is configured.

Important: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information.

-clean

(Optional) Specifies that the volume that is to become a secondary is clean. Any changes that are made on the secondary volume are ignored, and only changes made on the clean primary volume 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 be used only if all data changed on the secondary volumes while the consistency group was in the idling state matches the state of the primary volumes when the consistency group was stopped. Otherwise, relationships that are not consistent are reported as consistent. When this is completed there is no method to determine whether these volumes 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 **startrelationship** command starts a stand-alone relationship. The command fails if it is used to start a relationship that is part of a consistency group.

Note: You cannot start a relationship if the primary and secondary volumes are different sizes.

This command can be specified only 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.

Note: A command in idling state is rejected if any of the indicated secondary volumes is the target of an existing FlashCopy map.

If the FlashCopy mapping is active, the remote copy cannot be started.

If an existing remote copy relationship is stopped with **-access** specified but is restarted and the resultant secondary volume (depending on the choice of primary) is mapped to a host of type `hide_secondary`, that volume is not presented to the host. This is true even though it is mapped for configuration purposes. The mapped volumes are presented to the host if the:

- Host type is changed to a type other than `hide_secondary`
- Remote copy relationship is stopped and **-access** is specified
- Volume ceases to be a secondary volume because the remote copy relationship is being deleted or switched

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 situation can occur if input transactions occur on either the primary or secondary volumes since the ConsistentStopped or Idling state occurred. This situation occurs when the relationship is in either of these states:

- ConsistentStopped but not synchronized
- Idling but not synchronized

After restarting a relationship in either of these states, the data on the secondary volume is not usable for disaster recovery until the relationship becomes consistent.

A Global Mirror relationship with a `cycling_mode` of `multi` in either of these states does not require the **-force** parameter because a consistent secondary image is retained. However, if such a relationship is in idling state and written data has been received at the secondary volume, the **-force** flag is required because the secondary volume has a divergent image that cannot represent a consistent earlier state.

The **-force** parameter is not required if the relationship is in one of the following states:

- InconsistentStopped
- InconsistentCopying
- ConsistentSynchronized

However, the command does not fail if you specify the **-force** parameter.

You do not have to specify the **-force** parameter for relationships with configured secondary change volumes. If you specify **startcrrelationship** for an idling relationship, consistency protection is disabled if the secondary volume has been written to. This means that you must the **-force** parameter.

A Global Mirror relationship with a cycling mode of:

- *none* uses the non-cycling Global Mirror algorithm
- *multi* must:
 - Use a change volume that is configured at the primary volume (or the command fails)
 - Use a change volume that is configured at the secondary volume (or the command fails)
 - Perform multiple cycles of cycling

After creating a background copy the relationship remains in copying state, waiting for the remainder of the period time to expire before performing a new cycle. If the secondary change volume is unconfigured when the background copy completes, the relationship stops as if there is no cycle period.

Relationships that are active-active must have a state of idling to be started. (You must specify **-primary** to determine which of the master and auxiliary copies become the primary when starting an idling relationship.)

Use this command to:

- Restart the active-active relationship copy process and retain the historical disaster recovery copy that access has been granted to (which might have been used while the up-to-date copy was offline)
- Switch back to an up-to-date copy in the same state it was before **stopprrelationship -access** was specified. Any changes that are made to the historical copy are discarded

Remember: If you switch back to the up-to-date copy, you might have to take host actions to prepare for the volume data changing.

After specifying this command, if the secondary copy is not a historical copy of the primary relationship, it cannot be used for disaster recovery (and disaster recovery availability is restored after the copies are resynchronized). This situation can occur when:

- The new primary is the historical copy, which means the new secondary copy contains data that is from a later point in time than the data that the primary contains
- The secondary copy is the historical copy and was modified between the specification of **stopprrelationship -access** and **startcrrelationship -primary** command (Which means the secondary copy represents a divergent data image)

This command copies only the regions that are needed to resynchronize the two copies.

An invocation example

```
startcrrelationship rccopy1
```

The resulting output:

No feedback

stopprconsistgrp

Use the **stopprconsistgrp** command to stop the copy process for a Metro Mirror, Global Mirror, or active-active consistency group. This command can also be used to enable write access to the secondary volumes in the group if the group is in a consistent state.

Syntax

```
➤ stopprconsistgrp — [ -access ] [ rc_consist_group_id | rc_consist_group_name ] ➤
```

Parameters

-access

(Optional) Allows write access to consistent secondary volumes 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 volumes to secondary volumes.

Note: You cannot stop a consistency group using the **-access** parameter for a relationship if the primary and secondary volume are different sizes.

If the consistency group is in an inconsistent state, all copy operations stop and do not resume until you issue the **startprconsistgrp** command. When a consistency group is in a consistent state (*consistent_stopped*, *consistent_synchronized*, *consistent_copying*, or *consistent_disconnected*), you can issue the **access** parameter with the **stopprconsistgrp** command to enable write access to the secondary volumes within that group. For a consistency group in the *consistent_synchronized* state, this command causes a consistency freeze.

The *consistent_copying* state is a consistent state. A consistency group in this state transitions to *consistent_stopped* state if it receives a **stopprconsistgrp** command. Because the secondary change volume holds the consistent image, a stopped *consistent_copying* relationship might not have its secondary change volume deconfigured. This can be achieved by enabling access or completing synchronization so the secondary disk contains a consistent image. A relationship in *consistent_copying* or *consistent_stopped* accepts **stopprrelationship -access** transition to idling state.

The consistent image that is present on the change volume is made accessible at the secondary volume and after the command has completed the secondary volume can serve host read and write I/O.

If you specify **stopprconsistgrp -access** for a consistency group in a *consistent_copying* state the last consistent image on all the relationships in that group is restored. This process starts a FlashCopy mapping with the secondary change volume for the secondary volume in each relationship, which might cause the command to fail.

The relationship's data is from a different point in time than that of the consistency group if:

1. The consistency group is in a *consistent_copying* state
2. A relationship has been added to the group since the state became *consistent_copying*

Consequently, the relationship and consistency group are not be mutually consistent, and attempting to stop and enable access to the consistency group results in an error. To fix this, let a background copy complete so the consistency group becomes *consistent_synchronized*) or remove the inconsistent relationship from the consistency group before you enable access. If you stop the consistency group without the **-access** parameter, the consistency group becomes *consistent_stopped* but the secondary change volumes continue to retain a consistent image.

A FlashCopy background copy operation begins to migrate the data for the consistent image from the change volume to the secondary volume. While the background copy operation is in progress, the change volume for the secondary volume remains in use.

It might be necessary to process I/O before the reverse FlashCopy map can be triggered, causing the enable access command to time out. In this case, the relationship delays transitioning to idling until the reverse map starts and write access is available. Read access to the consistent data remains available.

To stop active-active consistency groups:

- **-access** is specified
- The relationship's state is `consistent_copying`
- The relationship's status is `primary_offline`

Specify `stopprconsistgrp -access` to obtain host read or write access to a volume in an active-active consistency group that contains an older but consistent image that might be needed in a disaster recovery scenario (the relationship has a `consistent_copying` state).

Any remote copy secondary volumes that are mapped to hosts of type `hide_secondary` are presented to the host if you specify **-access**. Paths to those volumes appear to the host, and a logical unit number (LUN) inventory changed unit attention is raised to report their availability.

Table 95 shows consistency group initial and final states:

Table 95. stopprconsistgrp 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 <code>idling</code> .
<code>consistent_synchronized</code>	<code>consistent_stopped</code>	If access is specified, the final state is <code>idling</code> . If access is not specified, the final state is <code>consistent_stopped</code> .
<code>consistent_copying</code>	<code>consistent_stopped</code>	If access is specified, the final state is <code>idling</code> . If access is not specified, the final state is <code>consistent_stopped</code> .
<code>idling</code>	<code>idling</code>	Remains in <code>idling</code> state whether access is specified or not.
<code>idling_disconnected</code>	<code>unchanged</code>	If specified without access , the relationship or group remains in <code>idling_disconnected</code> state. If the clustered systems 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 access flag.

Table 95. **stoprcconsistgrp** consistency group states (continued)

Initial state	Final state	Notes®
consistent_disconnected	consistent_stopped	The command is rejected if specified without access . If specified with access , the relationship or group moves to idling_disconnected.

An invocation example

```
stoprcconsistgrp rccopy1
```

The resulting output:

No feedback

stoprcrelationship

Use the **stoprcrelationship** command to stop 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 volume including for an active-active relationship.

Syntax

```

➤— stoprcrelationship — [ -access ] [ rc_rel_id | rc_rel_name ]

```

Parameters

-access

(Optional) Specifies that the system allow write access to a consistent secondary volume.

rc_rel_id | *rc_rel_name*

(Required) Specifies the ID or the name of the relationship to stop all processing for.

Description

The **stoprcrelationship** 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 volumes.

Note: You cannot stop a relationship using the **-access** parameter for a relationship if the primary and secondary volume are different sizes.

If the relationship is in an inconsistent state, any copy operation stops and does not resume until you issue a **startrcrelationship** 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, consistent_copying, or consistent_disconnected state – you can use the **access** parameter to enable write access to the secondary volume. Table 96 on page 577 provides consistency group initial and final states.

The consistent_copying state is a consistent state. A relationship in consistent_copying state transitions to consistent_stopped state when you specify **stoprcrelationship**. Because the secondary change volume holds the consistent image, a stopped consistent_copying relationship might not have its secondary change volume deconfigured. This can be achieved by enabling access or completing

synchronization so the secondary disk contains a consistent image. A relationship in `consistent_copying` or `consistent_stopped` accepts **stoprelationship -access** transition to idling state.

The consistent image that is present on the change volume is made accessible at the secondary volume. After the command completes, the secondary volume can serve host read and write I/O.

A FlashCopy background copy operation begins to migrate the data for the consistent image from the change volume to the secondary volume. While the background copy operation is in progress, the change volume for the secondary volume remains in use.

The enable access command can time out if there is I/O to process before the reverse FlashCopy map can be triggered. In this case, the relationship delays transitioning to idling until the reverse map starts and write access is available. Read access to the consistent data remains available.

To stop active-active relationships:

- **-access** is specified
- The relationship's state is `consistent_copying`
- The relationship's status is `primary_offline`

Specify **stoprelationship -access** to obtain host read or write access to a volume in an active-active relationship that contains an older but consistent image that might be needed in a disaster recovery scenario (the relationship has a `consistent_copying` state).

Any remote copy secondary volumes that are mapped to hosts of type `hide_secondary` are presented to the host if you specify **-access**. Paths to those volumes appear to the host, and a logical unit number (LUN) inventory changed unit attention is raised to report their availability.

When you enable read or write access for a `consistent_copying` relationship, specify **stoprelationship -access** to restore a consistent image on the secondary change volume that is using a FlashCopy mapping. (Depending on how long this operation lasts the CLI command might delay.) This process fails if either the secondary volume or secondary change volume are offline. If you stop the relationship without specifying the **-access** parameter, the relationship becomes `consistent_stopped` and the secondary change volume is unchanged.

To enable access to a secondary volume that is not consistent, specify **rmrelationship -force**.

Table 96. stoprelationship 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> .
<code>consistent_copying</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 96. stopprrelationship consistency group states (continued)

Initial state	Final state	Notes
idling	idling	Remains in idling state whether access is specified or not.
idling_disconnected	unchanged	If specified without access , the relationship or group remains in idling_disconnected state. If the clustered systems reconnect, the relationship or group is in either inconsistent_stopped or consistent_stopped 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 or group moves to idling_disconnected state.

An invocation example

```
stopprrelationship rccopy1
```

The resulting output:

No feedback

switchrcconsistgrp

Use the **switchrcconsistgrp** command to reverse the roles of the primary and secondary volumes 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

```
►► switchrcconsistgrp — — -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 volumes.

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 volumes in a consistency group, perhaps as part of a failover process that is associated with a disaster recovery event.

Note: You cannot switch a consistency group if the primary and secondary volumes are different sizes.

Write access to the former primary volumes is lost and write access to the new primary volumes 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 **switchrcconsistgrp** command to process correctly:

- ConsistentSynchronized
- ConsistentStopped and Synchronized
- Idling and Synchronized

Note: This command is rejected under any of the following conditions:

- You switch consistency group relationship so that the new secondary becomes the target volume of an active FlashCopy mapping.
- Any of the indicated secondary volumes (in the consistency group) are the target of an existing FlashCopy mapping.
- Using Global Mirroring with the multi cycling mode

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.

When the direction of the consistency group is changed, a volume that is a secondary volume in a remote copy relationship becomes a primary volume. In addition, a primary volume that is in a remote copy relationship becomes a secondary volume. If the resultant secondary volume is mapped to a host of type `hide_secondary`, it is no longer presented to that host. However, the mapping still exists for configuration purposes. If the volume that was a secondary volume before the switch is mapped to a host of type `hide_secondary`, it is presented to that host because it is no longer a secondary volume.

You cannot switch directions for an active-active consistency group.

An invocation example

```
switchrcconsistgrp -primary aux rccopy2
```

The resulting output:

No feedback

switchrcrelationship

Use the **switchrcrelationship** command to reverse the roles of primary and secondary volumes in a stand-alone Metro Mirror or Global Mirror relationship when that relationship is in a consistent state.

Syntax

```
➤ switchrcrelationship — — -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 **switchcrelationship** 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 volume in a relationship perhaps as part of a failover process during a disaster recovery event.

Note: You cannot switch a relationship if the primary and secondary volumes are different sizes.

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 **switchcrelationship** command to process correctly:

- ConsistentSynchronized
- ConsistentStopped and Synchronized
- Idling and Synchronized

Note: A command in idling state is rejected if any of the indicated secondary volumes is the target of an existing FlashCopy map.

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.

When the direction of the relationship is changed, a volume that is a secondary volume in a remote copy relationship becomes a primary volume, and a primary volume in a remote copy relationship becomes a secondary volume. If the resultant secondary volume is mapped to a host (type `hide_secondary`), it is no longer presented to that host. However, the mapping still exists for configuration purposes. If the volume that was a secondary volume before the switch is mapped to a host of type `hide_secondary`, it is presented to that host because it is no longer a secondary volume.

The **switchcrelationship** command is rejected if you use Global Mirroring with the **multi** cycling mode.

You cannot switch directions for an active-active relationship.

An invocation example

```
switchcrelationship -primary master rccopy2
```

The resulting output:

No feedback

Chapter 23. Migration commands

The migration commands enable you to work with migration options with the SAN Volume Controller.

lsmigrate

Use the **lsmigrate** command to display the progress of all current data migration operations.

Syntax

```
➤ lsmigrate — [ -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 are reported using this command. Easy Tier migrations are not included in the output.

An invocation example

```
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
```

migrateexts

Use the **migrateexts** command to migrate extents from one managed disk to another.

Syntax

```
►► migrateexts — — -source — [ source_mdisk_id — | — source_mdisk_name — ] — -target — [ target_mdisk_id — | — target_mdisk_name — ] — ►
► -exts — number_of_extents — — [ -threads — number_of_threads — ] [ -copy — id — ] — ►
► -vdisk — [ vdisk_id — | — vdisk_name — ] — ►
```

Parameters

- source** *source_mdisk_id* | *source_mdisk_name*
(Required) Specifies the MDisk on which the extents currently reside.
- target** *target_mdisk_id* | *target_mdisk_name*
(Required) Specifies the MDisk to migrate the extents to.
- exts** *number_of_extents*
(Required) Specifies the number of extents to migrate.
- threads** *number_of_threads*
(Optional) Specifies the number of threads to use while migrating these extents. You can specify 1 - 4 threads. The default number of threads is 4.
- copy** *id*
(Required if the specified volume has more than one copy) Specifies the volume copy that the extents belong to.
- vdisk** *vdisk_id* | *vdisk_name*
(Required) Specifies the volume that the extents belong to.

Description

This command migrates a given number of extents from the source volume and the managed disk that contains extents that are used to make up the volume. The target is a managed disk within the same storage pool.

If a large number of extents are being migrated, you can specify 1 - 4 threads. You can issue the **lsmigrate** command to check the progress of the migration.

The **migrateexts** command fails if there are insufficient free extents on the target managed disk. To avoid this problem, do not issue new commands that use extents until the extents migration is completed.

The **migrateexts** command fails if the target or source volume is offline, or if Easy Tier is active for the volume copy. Correct the offline condition before attempting to migrate the volume.

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 volume 1 created in it
- MDiskGrp 2 has volume 2 created in it
- MDiskGrp 3 has only one MDisk

With the previous configuration, the following migration operations are started:

- Migration 1 migrates volume 1 from MDiskGrp 1 to MDiskGrp 3, running with 4 threads.
- Migration 2 migrates volume 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.

Remember: This command cannot be used if the source MDisk is an SAS MDisk (which works in image mode only).

An invocation example

```
migrateexts -vdisk vdisk4 -source mdisk4 -exts
64 -target mdisk6 -threads 4
```

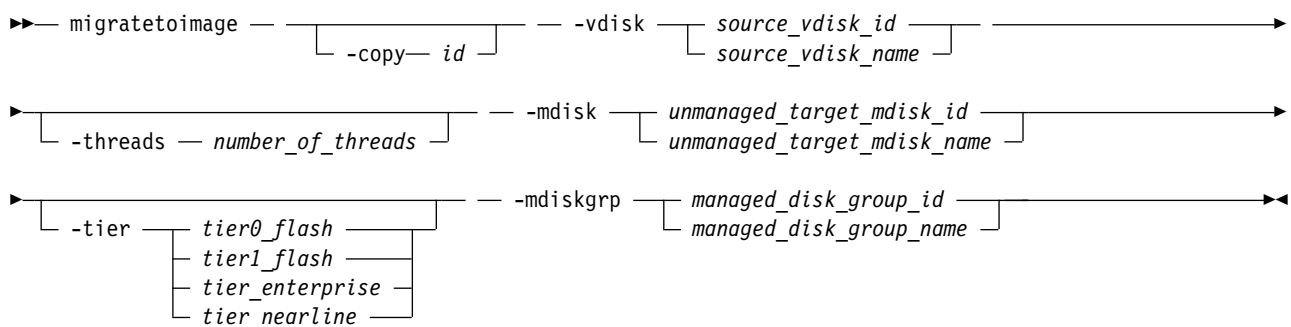
The resulting output:

No feedback

migratetoimage

Use the **migratetoimage** command to migrate data from a volume (image mode or managed mode) onto a new image mode volume copy. The target disk does not have to be in the same storage pool as the source disk.

Syntax



Parameters

-vdisk *source_vdisk_id* | *name*

(Required) Specifies the name or ID of the source volume to be migrated.

-copy *id*

(Required if the specified volume has more than one copy) Specifies the volume 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.

-mdisk *unmanaged_target_mdisk_id | name*

(Required) Specifies the name of the MDisk to which the data must be migrated. This disk must be unmanaged and large enough to contain the data of the disk that is being migrated.

-mdiskgrp *managed_disk_group_id | name*

(Required) Specifies the storage pool into which the MDisk must be placed, after the migration has completed.

-tier *tier0_flash | tier1_flash | tier_enterprise | tier_nearline*

(Optional) Specifies the tier of the MDisk being added.

tier0_flash

Specifies a tier0_flash hard disk drive or an external MDisk for the newly discovered or external volume.

tier1_flash

Specifies an tier1_flash (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

tier_enterprise

Specifies a tier_enterprise hard disk drive or an external MDisk for the newly discovered or external volume.

tier_nearline

Specifies a tier_nearline hard disk drive or an external MDisk for the newly discovered or external volume.

Description

The **migratetoimage** command migrates the data of a user-specified volume by consolidating its extents (which might reside on one or more MDisk) onto the extents of the target MDisk that you specify. After migration is complete, the volume is classified as an image type volume, and the corresponding mdisk is classified as an image mode MDisk.

Remember: This command cannot be used if the source volume copy is in a child pool or if the MDisk group that is specified is a child pool.

This command does not work if the volume is fast formatting.

Note: You cannot migrate a volume or volume image between storage pools if cloud snapshot is enabled on the volume.

The managed disk that is specified as the target must be in an unmanaged state at the time that the command is run. Running this command results in the inclusion of the MDisk into the user-specified storage pool.

The **migratetoimage** command fails if the target or source volume is offline. Correct the offline condition before attempting to migrate the volume.

Remember: This command cannot be used on a volume owned by a filesystem or if the source MDisk is an SAS MDisk (which works in image mode only).

If the volume (or volume copy) is a target of a FlashCopy mapping with a source volume in an active-active relationship the new managed disk group must be in the same site as the source volume. If the volume is in an active-active relationship the new managed disk group must be located in the

same site as the source volume. Additionally, the site information for the MDisk being added must be well-defined and match the site information for other MDisks in the storage pool.

Note: You cannot migrate data from a volume if the target volume's formatting attribute value is yes.

An encryption key cannot be used when migrating an image mode MDisk. To use encryption (when the MDisk has an encryption key), the MDisk must be self-encrypting.

An invocation example

The following example specifies that the user wants to migrate the data from `vdisk0` onto `mdisk5` and that the MDisk must be put into the storage pool `mdgrp2`.

```
migratetoimage -vdisk vdisk0 -mdisk mdisk5 -mdiskgrp mdgrp2 -tier tier_nearline
```

The resulting output:

No feedback

migratevdisk

Use the **migratevdisk** command to migrate an entire volume from one storage pool to another storage pool.

Syntax

```
►► migratevdisk — — -mdiskgrp [ mdisk_group_id | mdisk_group_name ] [ -threads — number_of_threads ]
► [ -copy — id ] — -vdisk [ vdisk_id | vdisk_name ] ◄◄
```

Parameters

-mdiskgrp *mdisk_group_id* | *mdisk_group_name*
(Required) Specifies the new storage pool ID or name.

-threads *number_of_threads*
(Optional) Specifies the number of threads to use during the migration of these extents. You can specify 1 - 4 threads. The default number of threads is 4.

-copy *id*
(Required if the specified volume has more than one copy) Specifies the volume copy to migrate.

-vdisk *vdisk_id* | *vdisk_name*
(Required) Specifies the volume ID or name to migrate in to a new storage pool.

Description

The **migratevdisk** command migrates the specified volume into a new storage pool; all the extents that make up the volume are migrated onto free extents in the new storage pool.

You can reassign a volume from a:

- Child pool to its parent pool
- Parent pool to one of its child pools
- Between the child pools in the same parent pool
- Between two parent pools

Note: You cannot migrate a volume between storage pools if cloud snapshot is enabled on the volume.

You can issue the **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 storage pool for the duration of the command. To avoid this problem, do not issue new commands that use extents until the volume migration is completed.

The **migratevdisk** command fails if the target volume or source volume is offline. Correct the offline condition before attempting to migrate the volume.

Remember:

- This command cannot be used on a volume owned by a filesystem.
- This command cannot be used if the source MDisk is an SAS MDisk (which works in image mode only).

If the volume (or volume copy) is a target of a FlashCopy mapping with a source volume in an active-active, relationship the new storage pool must be in the same site as the source volume. If the volume is in an active-active relationship the new storage pool must be located in the same site as the source volume.

When the volume is being migrated from a parent pool to another parent pool, the information is moved (unchanged), whether or not one or other is encrypted. The parent pool and the child pool cannot have an encryption key (or else the child pool would have failed during creation).

- A parent pool to parent pool migration is allowed in all cases.
- A parent pool to child pool migration is not allowed if child has encryption key.
- A child pool to parent pool or child pool is not allowed if either child pool has an encryption key.

An invocation example

```
migratevdisk -vdisk 4 -mdiskgrp Group0 -threads 2
```

The resulting output:

No feedback

Chapter 24. Service information commands

Service information commands are used to view the status of the hardware and to report hardware errors.

Note: You can find out the `panel_name` for each of your nodes by issuing `lservicenodes`.

host

Use the **host** command to change an Internet Protocol (IP) address to a host name or change a host name to an IP address.

Syntax

```
➤— sainfo host — — -ip_or_name — ip_address —————➤
                                   |
                                   host_name
```

Parameters

-ip *ip_address* | *host_name*

(Required) Specifies the host system IP address or host system name. The value for the IP address must be a standard IPv4 or IPv6 address. The value for the host name must be an alphanumeric string.

Description

This command changes an IP address to a host name or change a host name to an IP address.

This command handles Domain Name System (DNS) lookup and assists with DNS configuration.

Use this command to convert host names to IP addresses or IP address to host names.

An invocation example

```
sainfo host -ip_or_name varyd
```

The resulting output:

```
Host dummy not found: 3(NXDOMAIN)
```

An invocation example

```
sainfo host -ip_or_name compass.ssd.hursley.ibm.com
```

The resulting output:

```
compass.ssd.hursley.ibm.com has address 9.71.44.59
```

sainfo lsbootdrive

Use the **lsbootdrive** command to return the drive information of the internal boot drives for the given node. This command applies to SAN Volume Controller 2145-DH8 systems.

Syntax

►►—sainfo lsbootdrive— —panel_name—►►

Parameters

panel_name

(Optional) Identifies the node being used.

Description

The command displays information about internal boot drives for the specified node (if applicable).

Table 97 provides the possible values for attributes displayed as data in the output views.

Table 97. lsbootdrive attribute values

Attribute	Value
panel_id	Identifies the panel ID of the node that contains the boot drive. The value is a 7-character alphanumeric string.
node_id	Identifies the ID (in decimal format) of the node that contains the boot drive.
node_name	Identifies the name of the node that contains the boot drive.
can_sync	Indicates when synchronization is needed and is not prevented by another problem. The values are yes or no. Remember: This value can only be yes when status is out_of_sync. If the status is not out_of_sync the value is no.
slot_id	Identifies the ID (in decimal format) of the slot within the node.
booted	Indicates whether the node started from the specified drive. The values are yes, no.
status	Indicates the slot status. The values are: <ul style="list-style-type: none">• missing indicates that the slot should be occupied, but the software cannot detect a drive. It also indicates the serial number of the expected drive.• empty indicates that the slot is supposed to be empty and is empty.• unsupported indicates that the slot is supposed to be empty but is not.• failed indicates that the drive in the slot is not working.• uninitialized indicates that the drive is not formatted for the system.• wrong_node indicates that the drive is working but is not in the correct node. It also indicates the serial number of the node the drive must be in and might indicate the serial number of the drive that must be in the slot.• wrong_slot indicates that the drive is working correctly - it is in the right node but the wrong slot. It also indicates which drives belong in which slots.• out_of_sync indicates that the drive is working correctly but must be resynchronized. Make sure the value of can_sync is yes before specifying the command chbootdrive -sync.• online indicates that the drive in the slot is working correctly.• unknown indicates that the node is not an active member of the clustered system (system), and the state of the drive in that slot is unknown.
actual_drive_sn	Indicates the serial number of drive in the slot. The value is an alphanumeric string or blank if there is none
configured_drive_sn	Identifies the serial number of the drive that must be in the slot . The value is an alphanumeric string or blank if there is none.
actual_node_sn	Indicates the serial number of the node the drive (currently in the slot) belongs in. The value is an alphanumeric string or blank if there is none.

Table 97. *lsbootdrive* attribute values (continued)

Attribute	Value
identify	Indicates if <code>chnodebootdrive -identify</code> is specified. The values are: <ul style="list-style-type: none"> <code>on</code> indicates <code>chnodebootdrive -identify yes -slot</code> is specified <code>off</code> indicates <code>chnodebootdrive -identify no -slot</code> is specified <code>N/A</code> indicates that the drive slot cannot be identified.
FRU_part_number	Identifies the field-replaceable unit (FRU) part number of the drive. The value is a 7-character alphanumeric string or blank if there is none. The value is <code>N/A</code> if the drive is not supplied for your system.
FRU_identity	Identifies the 11S number that combines the manufacturing part number and serial number. The value is a 22-character alphanumeric string. The value is <code>N/A</code> if the drive is not supplied for your system.

: If the status is `out_of_sync` and `can_sync` is set to `no`, look for an indication to determine what must be fixed in order to make resynchronization possible.

An invocation example

```
sainfo lsbootdrive
```

The resulting output:

```
panel_id node_id node_name can_sync slot_id booted status actual_drive_sn configured_drive_sn actual_node_sn identify FRU_part_number FRU_identity
bfbfbf1 1 bfn1 no 1 yes online 1234567 1234567 bfbfbf1 off 90Y8878 11S49Y7427YXX6XK
bfbfbf1 1 bfn1 no 2 no missing 1234568
bfbfbf1 1 bfn1 no 3 empty
bfbfbf1 1 bfn1 no 4 empty
bfbfbf1 1 bfn1 no 5 empty
bfbfbf1 1 bfn1 no 6 empty
bfbfbf1 1 bfn1 no 7 empty
bfbfbf1 1 bfn1 no 8 no unsupported 12B0345
```

sainfolscmdstatus

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 fails if the *panel_name* ID is not in the list returned by **lsservicenodes**.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

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 example below.

An invocation example

```
sainfo lscmdstatus
```

The resulting output:

```
last_command satask cpfiles -prefix /dumps/test_cpf* -source 111825
last_command_status CMMVC8044I Command completed successfully.
T3_status
T3_status_data
cpfiles_status Complete
cpfiles_status_data Copied 2 of 2
snap_status Complete
snap_filename /dumps/snap.single.111896.130123.151657.tgz
installcanistersoftware_status
```

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

(Optional) The name of the panel. The command will fail if the *panel_name* ID is not in the list returned by the **lsservicenodes** command.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

-prefix path

(Optional) 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/cloud
- /dumps/configs
- /dumps/drive
- /dumps/easytier
- /dumps/elog
- /dumps/enclosure
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/mdisk
- /dumps/syslogs
- /home/admin/update

Description

This command displays a list of the files on the node that you want to retrieve using the **satask cpfiles** command.

An invocation example to list the files in the /dumps directory

```
sainfo lsfiles -prefix /dumps
```

The resulting output:

```
filename
sublun.trc.old
sublun.trc
100050.trc.old
eccore.100050.100305.183051
eccore.100050.100305.183052
ethernet.100050.trc
100050.trc
```

An invocation example to list the files in the /dumps/easytier directory

```
sainfo lsfiles -prefix /dumps/easytier/ 01-1
```

The resulting output:

```
filename
dpa_heat.78RE5LV-1.150705.074636.data
dpa_log_78RE5LV-1_20150707062320_00000000.xml.gz
```

lshardware

Use the **lshardware** command to view the configured and actual hardware configuration of a node in the clustered system (system).

Syntax

```
➤ sainfo — lshardware — [ -delim — delimiter ] [ panel_name ] ➤
```

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.

panel_name

(Optional) The node panel name.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

Description

When the node is in a service state, use this command to view the current hardware configuration. Table 98 on page 592 provides the possible values that are applicable to the attributes that are displayed as data in the output views.

Table 98. *lshardware* attribute values

Attribute	Value
panel_name	Indicates the node panel name.
node_id	Indicates the node unique ID; blank if not in a system.
node_name	Indicates the node name; blank if not in a system.
node_status	Indicates the node status.
hardware	Indicates the hardware model, such as CG8.
actual_different	Indicates if the node hardware is different than the configured hardware.
actual_valid	Indicates if the node hardware is valid.
memory_configured	Indicates the configured amount of memory (in GB).
memory_actual	Indicates the currently installed amount of memory (in GB).
memory_valid	Indicates if the actual memory is a valid configuration.
cpu_count	Indicates the maximum number of CPUs for the node.
cpu_socket	Indicates the ID of socket to which the CPU fields refer.
cpu_configured	Indicates the configured CPU for this socket.
cpu_actual	Indicates the currently installed CPU in this socket.
cpu_valid	Indicates if the currently installed CPU is a valid configuration.
adapter_count	Indicates the maximum number of adapters for the node (differs by node type).
adapter_location	Indicates the location of this adapter.
adapter_configured	Indicates the configured adapter for this location.
adapter_actual	Indicates the currently installed adapter for this location.
adapter_valid	Indicates if the adapter in this location is valid.
ports_different	Indicates whether adapter ports can support more functions.

An invocation example for a 300 node

```
sainfo lshardware
```

The resulting output:

```
panel_name,123456
node_id,
node_name
status,service
hardware,300
actual_different,no
actual_valid,yes
memory_configured,8
memory_actual,8
memory_valid,yes
cpu_count,1
cpu_socket,1
cpu_configured,4 core Intel(R) Xeon(R) CPU E3110 @ 3.0GHz
cpu_actual,4 core Intel(R) Xeon(R) CPU E3110 @ 3.0GHz
cpu_valid,yes
cpu_socket,
cpu_configured,
cpu_actual,
cpu_valid,
adapter_count,6
adapter_location,0
adapter_configured,1Gb/s Ethernet adapter
adapter_actual,1Gb/s Ethernet adapter
```

```

adapter_valid,yes
adapter_location,0
adapter_configured,1Gb/s Ethernet adapter
adapter_actual,1Gb/s Ethernet adapter
adapter_valid,yes
adapter_location,0
adapter_configured,Four port 8Gb/s FC adapter card
adapter_actual,Four port 8Gb/s FC adapter card
adapter_valid,yes
adapter_location,0
adapter_configured,High-speed SAS adapter
adapter_actual,High-speed SAS adapter
adapter_valid,yes
adapter_location,0
adapter_configured,Midplane bus adapter
adapter_actual,Midplane bus adapter
adapter_valid,yes
adapter_location,1
adapter_configured,Two port 10Gb/s ethernet adapter
adapter_actual,Two port 10Gb/s ethernet adapter
adapter_valid,yes
ports_different,no

```

Isservicenodes

Use the **Isservicenodes** command to display a list of all the nodes that can be serviced using the service assistant CLI.

Syntax

►► — **sainfo** — — **Isservicenodes** —————►►

Parameters

None

Description

This command displays a list of all the nodes that can be serviced using the service assistant CLI. This list includes nodes that at a code level of at least 6.2.0, are visible on the fabric, and are one of the following:

- The partner node in a control enclosure to the node that is running the command.
- In the same clustered system as the node running the command.
- In candidate state.
- Not in a clustered system and in service state.
- Not in an enclosure with a stored clustered system ID (which is not the clustered system ID of the local node).

Nodes not clustered with the local node will not be shown unless they are the partner node. Table 99 shows possible outputs.

Table 99. 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.

Table 99. *Isservicenodes outputs (continued)*

Attribute	Value
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 system as the local node. Candidate: the node is not part of the system.
node_status	The values are: <ul style="list-style-type: none"> Active indicates the node is part of a system and can perform I/O. Service indicates the node is in service, standby, or node rescue. Candidate indicates the node is not part of a system. Starting indicates the node is part of a system and is attempting to join the system, and cannot perform I/O.
error_data	Outstanding error and error data indicated by priority.

An invocation example

```
sainfo lsservicenodes
```

The resulting output:

```
panel_name cluster_id      cluster_name      node_id node_name relation node_status error_data
01-1       0000020073C0A0D4 Cluster_9.180.28.82 1      node1    local    Active
01-2       0000020073C0A0D4 Cluster_9.180.28.82 2      node2    partner  Active
```

Isservicerecommendation

Use the **lsservicerecommendation** command to determine what actions must be completed when servicing a node.

Syntax

```
➤— sainfo — — lsservicerecommendation —————▶
                                   |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 **lsservicenodes** 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 **lsservicenodes**.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

Description

This command enables you to determine what actions need to be completed when servicing a node.

An invocation example

Example 1 for `service_action`:

```
sainfo lsservicerecommendation
```

The resulting output:

Use fabric tools to diagnose and correct Fibre Channel fabric problem.

An invocation example

Example 2 for `service_action`:

```
sainfo lsservicerecommendation
```

The resulting output:

No service action required, use console to manage node.

lsservicestatus

Use the **lsservicestatus** command to display the current status of a node.

Syntax

```
➤— sainfo — — lsservicestatus ————— panel_name —————➤
```

Parameters

panel_name

(Optional) If a *panel_name* is provided, the service recommendation for the local node is returned. If a *panel_name* from the list returned by **lsservicenodes** is specified, then the service recommendation for that node is returned. The command fails if the *panel_name* ID is not in the list returned by **lsservicenodes**. This output is returned as the node status on all Universal Serial Bus (USB) flash drive commands.

Note: For 2076 nodes the *panel name* is the value of the enclosure ID and canister ID or the enclosure serial number and canister location.

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 clustered system (system), to obtain the vital product data (VPD) and error status.

Table 57 on page 315 shows possible outputs.

Table 100. lsservicestatus output

Attribute	Value
panel_name	The front panel name, enclosure IDs, or canister IDs that identify the node.
console_ip	An Internet Protocol (IP) Version 4 or 6 address Note: This field might be blank if the node is not present in a system.
has_nas_key	yes no Note: This field might be blank if the node is not present in a system.
cluster_id	Specifies the ID of a system.

Table 100. Isservicestatus output (continued)

Attribute	Value
cluster_name	Specifies the name of a system. When you use this parameter, the detailed view of the specific system 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	4F2 CF8 CG8 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 version of code for the node.
node_sw_build	The build string version of code on the node.
cluster_sw_build	The CSM build that the system is running.
node_error_count	The number of node errors.
node_error_data	The type of node errors.
FC_port_count	The number of FC ports.
FC_port_id	The beginning of repeating fields for each FC port; the whole set of fields indicated is repeated for each port.
port_status	This should match the port on the front panel, enclosure, or canister.
port_speed	This should match the port speed on the front panel, enclosure, or canister.
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	10Mbps 100Mbps 1Gbps 10Gbps full half
MAC	A single MAC address.
vnport_count	Number of VN ports created on top of each physical FCoE port.
vnport_id	The VN port ID.
vnport_wwpn	The WWPN assigned to the VN port.
vnport_FCF_mac	The MAC address for the FCF to which the VN port is connected.
vnport_vlanid	The VLAN ID used by the VN port. The value is blank for FC ports.
product_mtn	The machine type and model.
product_serial	The node serial number.
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
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.
Battery_count	The number of expected batteries (two).
Battery_id	Specifies the ID of the slot the battery is in.
Battery_status	The status is missing, failed, charging, or active.

Table 100. *Isservicestatus output (continued)*

Attribute	Value
Battery_FRU_part	The FRU part number of the battery.
Battery_part_identity	The 11S FRU identity of the battery (includes the serial number).
Battery_fault_led	The fault light-emitting diode (LED) status.
Battery_charging_status	The battery charging status.
Battery_cycle_count	The number of charge or discharge cycles performed by the battery.
Battery_powered_on_hours	The number of hours the battery is in a powered node.
Battery_last_recondition	A system timestamp indicating the last successful gas gauge calibration.
Battery_midplane_FRU_part	The FRU part number of the battery midplane
Battery_midplane_part_identity	The 11S FRU identity of the battery midplane (which includes the serial number).
Battery_midplane_FW_version	The firmware version running on the battery midplane.
Battery_power_cable_FRU_part	The FRU part number of the battery power cable.
Battery_power_sense_cable_FRU_part	The FRU part number of the battery midplane power sense cable.
Battery_comms_cable_FRU_part	The FRU part number of the battery midplane communication cable.
Battery_EPOW_cable_FRU_part	The FRU part number of the battery midplane EPOW cable.
PSU_count	N/A
PSU_id	N/A
PSU_status	N/A
local_fc_port_mask	Indicates the FC I/O ports that a system can use for node-to-node communications on a local system if those FC I/O ports exist on a node. The value is 64 binary bits.
partner_fc_port_mask	Indicates the FC I/O ports that a system can use for system-to-system communications on a partnered system if those FC I/O ports exist on a node. The value is 64 binary bits.
cluster_topology	Indicates the system topology (set using the chsystem command).
site_id	Indicates the site node value.
site_name	Indicates the site name.
identify_LED	Indicates the node or node canister identify LED state (on, off, or blank).
password_reset_enabled	Indicates if the superuser password reset is enabled (yes or no).

An invocation example

```
sainfo lsservicestatus
```

The resulting output:

```
panel_name 150434
cluster_id 000002006ee1445e
cluster_name Cluster_192.168.8.241
cluster_status Active
cluster_ip_count 2
cluster_port 1
cluster_ip 192.168.8.241
cluster_gw 192.168.8.1
cluster_mask 255.255.255.0
cluster_ip_6
cluster_gw_6
cluster_prefix_6
cluster_port 2
cluster_ip
cluster_gw
cluster_mask
cluster_ip_6
cluster_gw_6
cluster_prefix_6
node_id 1
node_name node1
node_status Active
config_node Yes
hardware CF8

service_IP_address
service_gateway
service_subnet_mask
service_IP_address_6
service_gateway_6
```

```

service_prefix_6
service_IP_mode dhcpfallback
node_sw_version 6.4.0.0
node_sw_build 64.8.1205180000
cluster_sw_build 64.8.1205180000
node_error_count 0
fc_ports 4
port_id 1
port_status Active
port_speed 8Gb
port_WWPN 500507680140a22f
SFP_type Short-wave
port_id 2
port_status Active
port_speed 8Gb
port_WWPN 500507680130a22f
SFP_type Short-wave
port_id 3
port_status Active
port_speed 8Gb
port_WWPN 500507680110a22f
SFP_type Short-wave
port_id 4
port_status Active
port_speed 8Gb
port_WWPN 500507680120a22f
SFP_type Short-wave
ethernet_ports 4
ethernet_port_id 1
port_status Link Online
port_speed 1Gb/s - Full
MAC 00:21:5e:db:30:38
vnport_count 0
ethernet_port_id 2
port_status Not Configured
port_speed
MAC 00:21:5e:db:30:3a
vnport_count 0
ethernet_port_id 3
port_status Not Configured
port_speed 10Gb/s - Full
MAC 00:00:c9:bc:6f:22
vnport_count 0
ethernet_port_id 4
port_status Not Configured
port_speed 10Gb/s - Full
MAC 00:00:c9:bc:6f:20
vnport_count 0
product_mtm 2145-CF8

product_serial 75HAXYA
time_to_charge 0
Battery_charging 0
Battery_count 2
Battery_id 1
Battery_status active
Battery_FRU_part 12Z9876
Battery_part_identity 11S98Z1234YM11BG123456
Battery_fault_led OFF
Battery_charging_Status charged
Battery_cycle_count 5
Battery_powered_on_hours 12345
Battery_last_recondition 130629123456
Battery_id 2
Battery_status failed
Battery_FRU_part 12Z9876
Battery_part_identity 11S98Z1234YM11BG234567

```


Battery_fault_led ON
Battery_charging_Status charged
Battery_cycle_count 5
Battery_power_on_hours 12345
Battery_last_recondition 130702123400
Battery_midplane_FRU_part 12Z9880
Battery_power_cable_FRU_part 12Z9881
Battery_power_sense_cable_FRU_part 12Z9882
Battery_comms_cable_FRU_part 12Z9883
Battery_EPOW_cable_FRU_part 12Z9884

dump_name 150434
node_WWNN 500507680100a22f
disk_WWNN_suffix 0A22F
panel_WWNN_suffix 0A22F
UPS_serial_number
UPS_status
enclosure_WWNN_1
enclosure_WWNN_2
node_part_identity
node_FRU_part
enclosure_identity
PSU_count
PSU_id
PSU_status
PSU_id
PSU_status

node_location_copy
node_product_mtm_copy
node_product_serial_copy
node_WWNN_1_copy
node_WWNN_2_copy
latest_cluster_id
next_cluster_id
console_IP 192.168.8.241:443
has_nas_key no
fc_io_ports 6
fc_io_port_id 1
fc_io_port_WWPN 500507680140a22f
fc_io_port_switch_WWPN 200000051e630f9a
fc_io_port_state Active
fc_io_port_FCF_MAC N/A
fc_io_port_vlanid N/A
fc_io_port_type FC
fc_io_port_type_port_id 1
fc_io_port_id 2
fc_io_port_WWPN 500507680130a22f
fc_io_port_switch_WWPN 200400051e630f9a
fc_io_port_state Active
fc_io_port_FCF_MAC N/A
fc_io_port_vlanid N/A
fc_io_port_type FC
fc_io_port_type_port_id 2
fc_io_port_id 3
fc_io_port_WWPN 500507680110a22f
fc_io_port_switch_WWPN 200000051e7ded49
fc_io_port_state Active
fc_io_port_FCF_MAC N/A
fc_io_port_vlanid N/A
fc_io_port_type FC
fc_io_port_type_port_id 3
fc_io_port_id 4
fc_io_port_WWPN 500507680120a22f
fc_io_port_switch_WWPN 200400051e7ded49
fc_io_port_state Active
fc_io_port_FCF_MAC N/A

[illegible]

traceroute

Use the **tracert** command to print wide area network (WAN) route packets to a specified host.

Syntax

►► sainfo traceroute — — -ip_or_name *ip_address*
host name

Parameters

-ip *ip_address* | *host_name*
(Required) Specifies the host system IP address or host system name that you want to print WAN route packets to. The value for the IP address must be a standard IPv4 or IPv6 address. The value for the host name must be an alphanumeric string.

Description

This command prints WAN print route packets to a specified host.

An invocation example

```
sainfo traceroute -ip or name compass.ssd.hursley.ibm.com
```

The resulting output:

```
tracert to compass.ssd.hursley.ibm.com (9.71.44.59), 30 hops max, 60 byte packets
 1  9.71.45.4 (9.71.45.4)  0.283 ms  0.337 ms  0.397 ms
 2  compass.ssd.hursley.ibm.com (9.71.44.59)  0.124 ms  0.124 ms  0.113 ms
```

Chapter 25. Service mode commands (Discontinued)

The service mode commands are discontinued.

applysoftware (Discontinued)

Attention: The `svcservicemodetask applysoftware` command is discontinued. Use the `satask installsoftware` command instead.

Discontinued.

svcservicemodetask cleardumps (Discontinued)

Attention: The `svcservicemodetask cleardumps` command is discontinued. Use the `cleardumps` command instead.

Discontinued.

svcservicemodetask dumperrlog (Discontinued)

Attention: The `svcservicemodetask dumperrlog` command is discontinued. Use the `dumperrlog` command instead.

Discontinued.

exit (Discontinued)

Attention: The `svcservicemodetask exit` command is discontinued. Use the `satask stopservice` command instead.

Chapter 26. Service mode information commands (Discontinued)

The service mode information commands are discontinued.

ls2145dumps (Discontinued)

The `svcservicemodeinfo ls2145dumps` command is discontinued. Use the `lsdumps` command to display a list of files in a particular dumps directory.

lscimomdumps (Discontinued)

The `svcservicemodeinfo lscimomdumps` command is discontinued. Use the `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)

The `svcservicemodeinfo lserrlogdumps` command is discontinued. Use the `lsdumps` command to display a list of files in a particular dumps directory.

lsfeaturedumps (Discontinued)

The `svcservicemodeinfo lsfeaturedumps` command is discontinued. Use the `lsdumps` command to display a list of files in a particular dumps directory.

lsiostatsdumps (Discontinued)

The `svcservicemodeinfo lsiostatsdumps` command is discontinued. Use the `lsdumps` command to display a list of files in a particular dumps directory.

lsiotracedumps (Discontinued)

The `svcservicemodeinfo lsiotracedumps` command is discontinued. Use the `lsdumps` command to display a list of files in a particular dumps directory.

lsmdiskdumps (Discontinued)

The `svcservicemodeinfo lsmdiskdumps` command is discontinued. Use the `lsdumps` command to display a list of files in a particular dumps directory.

lssoftwaredumps (Discontinued)

The `svcservicemodeinfo lssoftwaredumps` command is discontinued. Use the `lsdumps` command to display a list of files in a particular dumps directory.

Chapter 27. Service task commands

Service task command information enables you to service the node hardware.

Note: You can find out the `panel_name` for each of your nodes by issuing **`sainfo lsservicenodes`**.

satask chbootdrive

Use the **chbootdrive** command to synchronize a broken drive or field-replaceable unit (FRU) replacement drive. This command applies to SAN Volume Controller 2145-DH8 systems.

Syntax

```
▶▶ satask chbootdrive — — -sync — — panel_name —▶▶
```

Parameters

-sync

(Required) Specifies synchronization for:

- New drives
- Drives marked out of synchronization
- Drives from another node

panel_name

(Optional) Specifies the panel ID of the node on which to operate.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Description

The command synchronizes a broken or FRU replacement drive.

Remember: This command can only be run in service mode (to keep the drive boots synchronized), and can only be used on systems equivalent to SAN Volume Controller 2145-DH8. Otherwise, use the **chnodebootdrive** command.

An invocation example

```
satask chbootdrive -sync
```


The resulting output:

No feedback

chnodeled

Use the **chnodeled** command to turn the identification light-emitting diode (LED) on or off for the specified node or control canister.

Syntax



Parameters

-on | -off

(Required) Turns the location LED for the specified node or control canister on or off.

-battery | -bootdrive slot

(Optional) Turns the LED on the battery or bootdrive on or off.

Looking at the front of the node:

- The value for the drive on the left side is 1 for slot 1. The value for the drive on the right side is 2 for slot 2.
- The battery on the left goes in the first battery slot (1). The battery on the right goes in the second battery slot (2).

If you specify **-battery** or **-bootdrive**, you must specify *slot*. If you do not specify **-battery** or **-bootdrive**, the node automatically sets the light-emitting diode (LED) identity for the control canister.

panel_name

(Optional) Specifies a unique panel name used to apply the command to a valid node on the fabric.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

Description

This command turns the canister location LED on or off.

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.

An invocation example to turn on the LED to identify the battery in slot 2 on node KP2812

```
satask chnodeled -on -battery 2 KP2812
```

The resulting output:

No feedback

An invocation example to turn on the LED to identify the battery on node 2

```
satask chnodeled -on -battery 2
```

The resulting output:

No feedback

An invocation example to allow the node to identify the LED

```
satask chnodeled -on
```

The resulting output:

No feedback

An invocation example to turn on the LED for canister 1 in control enclosure 02

```
satask chnodeled -on 02-1
```

The resulting output:

No feedback

chserviceip

Use the **chserviceip** command to set the service 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 — [-dhcp] [-dhcp_6] [-dhcpfallback] [-resetpassword] [-techport enable|disable] [-force] [panel_name]
```

Parameters

panel_name

(Optional) Identifies the node being serviced.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

-serviceip

(Optional) 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

(Optional) The IPv4 gateway for the service assistant. If **-gw** is specified, **-mask** must be specified.

-mask

(Optional) The IPv4 subnet for the service assistant. If **-mask** is specified, **-gw** must be specified.

-serviceip_6

(Optional) The Internet Protocol Version 6 (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

(Optional) The IPv6 gateway for the service assistant. If **-gw_6** is specified, **-prefix_6** must be specified.

-prefix_6

(Optional) The IPv6 prefix for the service assistant. If **-prefix_6** is specified, **-gw_6** must be specified.

-default

(Optional) Resets to the default IPv4 address.

-dhcp

(Optional) Attempts to obtain an IPv4 address from Dynamic Host Configuration Protocol (DHCP).

-dhcp_6

(Optional) Attempts to obtain an IPv6 address from DHCP.

-default

(Optional) Resets the IPv4 service address to the default address.

-resetpassword

(Optional) 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 clustered system (system) then the system 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 are preserved.

When **-dhcpfallback** is specified, the current service interface is restarted and the new service IPv4 address is established using DHCP. If the DHCP request fails, the service IP address is set statically based on the node's physical location.

The location is based on the:

- Enclosure location within the chassis
- Node slot within the enclosure

Do not use the **-dhcpfallback** parameter for IPv6. These flags allocate a new address because the command causes the service interface to restart.

Consequently, you can configure both an IPv4 and an IPv6 address concurrently. Setting the IPv4 address does not change the IPv6 setting, and setting the IPv6 address will not change the IPv4 setting. It is possible to clear any values set by setting the IPv4 address to 0.0.0.0 or leaving the IPv6 value empty.

Use the **chserviceip** command to:

- Clear the IPv4 service IP address:
`satask chserviceip -serviceip 0.0.0.0 -gw 0.0.0.0 -mask 0.0.0.0`
- Clear the IPv6 service IP address:
`satask chserviceip -serviceip_6 0::0 -gw_6 0::0 -prefix_6 64`

Remember:

- If **-gw** is specified, **-mask** must also be specified.
- If **-gw_6** is specified, **-prefix_6** must also be specified.

An invocation example

```
satask chserviceip
```

The resulting output:

No feedback

An invocation example using specific **-serviceip**, **-gw**, and **-mask** parameters

```
satask chserviceip -serviceip 1.2.3.4 -gw 1.2.3.1 -mask 255.255.255.0
```

The resulting output:

No feedback

An invocation example enabling the technician port on a node

```
satask chserviceip -techport enable -force
```

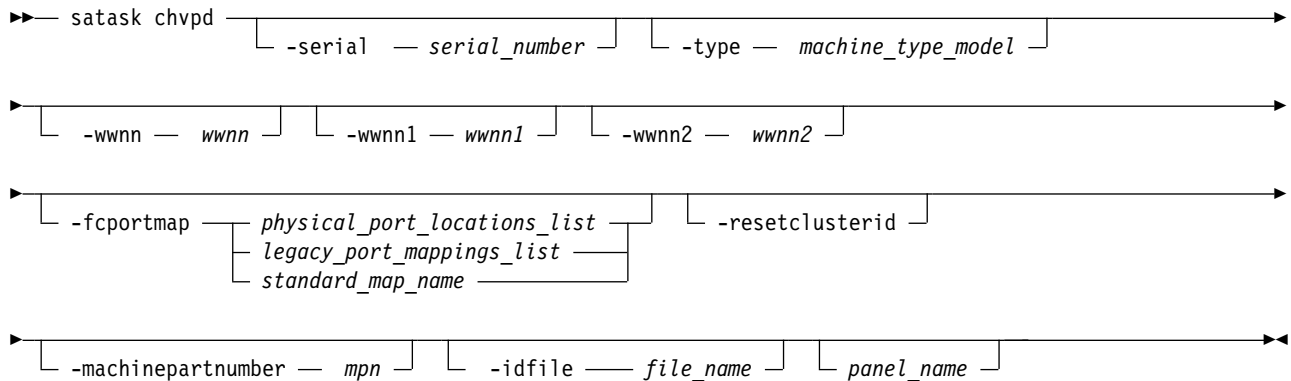
The resulting output:

No feedback

satask chvpd

Use the **chvpd** command to set vital product data (VPD) such as serial number and machine type.

Syntax



Parameters

-serial *serial_number*

(Optional) Specifies the serial number for system board or enclosure.

-type *machine_type_model*

(Optional) Specifies the machine model type.

-wwnn *wwnn*

(Optional) Specifies the worldwide node name (WWNN).

wwnn1 *wwnn1*

(Optional) Specifies the WWNN of canister 1.

-wwnn2 *wwnn2*

(Optional) Specifies the WWNN of canister 2.

-fcportmap *physical_port_locations_list* | *legacy_port_mappings_list* | *standard_map_name*

(Optional) Specifies the FC I/O port location on the node that is mapped to a slot or port.

- Specify `-fcportmap physical_port_locations_list` to associate physical ports with an FC I/O port ID and to use the port WWPNNs that is defined for the FC I/O port ID. The physical port location is made up of 2 digits, which represent the slot number (first digit) and port number (second digit).

Note: This value is a comma-separated list of physical port locations.

- Specify `-fcportmap legacy_port_mappings_list` to associate physical ports with an FC I/O port ID and to use the WWPN value that is defined for the port location of the replaced node. This means that the WWPN for the port on the new node is the same as the mapped port on the replaced node. A legacy port mapping maps a physical port location on the new system to a physical port location on the replaced system.

Note: This value is a comma-separated list of legacy port mappings.

- Specify `-fcportmap standard_map_name` to specify a standard port mapping.

Note: The FC I/O port ID of each port is defined by its position in the list.

-machinepartnumber *mpn*

(Optional) Specifies the machine part number.

-resetclusterid

(Optional) Resets the clustered system (system) ID to 0.

-idfile *file_name*

(Optional) Specifies the unique node ID file that allows each supported server to become a unique system node.

panel_name

(Optional) Identifies the node that is used.

Description

The command sets vital product data (VPD), including the serial number, WWPN, machine type, and system ID.

For port location management, the first location represents FC I/O port 1, which is the bit position furthest to the right in the port mask value. Any number of locations can be entered (including for adapters that are not yet installed).

The method for assigning WWPNs is different for each node and is selected automatically based on the WWNN of the node. You must set the mapping in the new node before adding it to the existing clustered system.

Note: When you change the mapping the node reboots to apply the new setting.

If you specify a WWNN for a legacy node, the port location from the original node must also be specified so the original WWPN can be assigned. If you use port masking, you must specify the ports in ascending slot or port order (from the original node).

A null mapping (specify `-fcportmap 00`) can be entered, which uses the default mapping. Do not change the port mapping for a node if that node is a member of a clustered system.

Remember: When you specify **chvdp**, one or more nodes might experience a reset or a reboot. For example, these parameters might cause a reboot:

- **-fcportmap** resets a node or one node canister
- **-idfile** resets a node
- **-serial** reboots both nodes or node canisters
- **-type** reboots both nodes or node canisters
- **-wwnn** resets a node or one node canister

An invocation example

```
satask chvdp -wwnn 1111111111111111
```

The resulting output:

No feedback

An invocation example

```
satask chvdpd -wwnn1 1111111111111111 -wwnn2 2222222222222222
```

The resulting output:

No feedback

An invocation example

```
satask chvdpd -serial 8675309
```

The resulting output:

No feedback

An invocation example

```
satask chvdpd -wwnn 500507680f000123 -fcportmap 31,32,33,34,41,42,43,44,61,62,63,64,71,72,73,74
```

The resulting output:

No feedback

An invocation example

```
satask chvdpd -wwnn 5005076801000456 -fcportmap 31-11,32-12,33-13,34-14,41-21,42-22,43-23,44-24,61-31,62-32,63-33,64-34,71-51,72-52
```

The resulting output:

No feedback

An invocation example

```
satask chvdpd -wwnn 5005076801000456 -fcportmap default
```

The resulting output:

No feedback

chwwnn

Use the **chwwnn** command to modify the node World Wide Node Name (WWNN). (This command applies to SAN Volume Controller 2145-CG8 and older nodes.)

Syntax

```
▶▶ satask — chwwnn — — -wwnnsuffix — wwnn_suffix — ┌──────────────────┐ ───────────────────▶▶
                                     └── panel_name ─┘
```

Parameters

-wwnnsuffix

(Required) The suffix to be used for node *wwnn*.

panel_name

(Optional) Identifies the node being serviced.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

Description

This command modifies the WWNN. Use the **lsservicestatus** command to view suggested WWNNs.

An invocation example

```
chwwnn -wwnnsuffix 000cc
```

The resulting output:

No feedback

cpfiles

Use the **cpfiles** command to copy files from another node.

Syntax

```
▶▶—satask— cpfiles—-prefix— $\left[ \begin{array}{l} \text{directory} \\ \text{file\_filter} \end{array} \right]$  —source —source_panel_name————→  
▶ $\left[ \text{target\_panel\_name} \right]$ ————→▶▶
```

Parameters

satask

System Administrator task; service commands that are only used in specific circumstances.

-prefix *directory* | *file_filter*

(Required) 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/cloud
- /dumps/configs
- /dumps/drive
- /dumps/easytier
- /dumps/elog
- (Storwize V7000)/dumps/enclosure
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/mdisk
- /dumps/syslogs
- /home/admin/update

Note:

- You can also specify a file filter. For example, if you specify /dumps/elog/*.txt, all files in the /dumps/elog 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/elog/*.txt"

-source *source_panel_name*

(Required) Identifies the source node files will be copied from.

target panel name

(Optional) Identifies the node that files are copied to. If no panel name is provided, the files are copied to the local node.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Description

This command copies files from another node. You can monitor the progress of the copy using the **sainfo1scmdstatus** command.

An invocation example to put configuration information from canister 1 in enclosure 2

```
satask cpfiles -prefix /dumps/configs -source 02-1
```

The resulting output:

No feedback

An invocation example to put Easy Tier information from canister 2 in enclosure 1

```
cpfiles -prefix /dumps/easytier/ -source 01-1 01-2
```

The resulting output:

No feedback

dumpinternallog (Discontinued)

The **dumpinternallog** command has been discontinued.

installsoftware

Use the **installsoftware** command to install a specific system code (code) package.

Syntax

►► `satask` — `installsoftware` — — `-file` *filename* — `-ignore` `-pacedccu` *panel name*

Parameters

-file *filename*

(Required) The file name of code installation package.

Note: The argument to **-file** must be present on the local node; the argument is automatically copied to the target panel_name.

-ignore

(Optional) Overrides prerequisite checking and forces installation of the code.

-pacedccu

(Optional) Causes the node to initiate a paced Concurrent Code Update (in which you define when the node begins its update) instead of a normal Concurrent Code Update (in which each node in the clustered system automatically updates in sequence).

panel_name

(Optional) Identifies the node being serviced.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Description

This command installs a specific code package.

An invocation example

```
satask installsoftware -file install_pkg.gpg nodeB_panel_name
```

The resulting output:

No feedback

leavecluster

Use the **leavecluster** command to remove clustered system (system) state data, location information, and other history from a node.

Syntax

```
▶▶ satask — leavecluster — — -force — — panel_name —▶▶
```

Parameters

-force

(Required) 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. The default is the node on which the command is entered.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Description

Use this command to remove system state data, location information, and other history from a node.

An invocation example

```
satask leavecluster -force 78G00F3-2 /* this forces the node with panel_name=78G00F3-2 out of the clustered system */
```

The resulting output:

No feedback

An invocation example

```
satask leavecluster -force /* this forces the node on which the command is entered out of the clustered system*/
```

The resulting output:

No feedback

An invocation example

`satask leavecluster -force /* this forces the node on which the command is entered out of the clustered system*/`

The resulting output:

No feedback

An invocation example

`satask leavecluster`

The resulting output:

CMMVC8034E A compulsory parameter is missing.

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

satask

System Administrator task; service commands that are only used in specific circumstances.

-rebuildcluster

(Required) Creates a cluster from the metadata found in `/dumps/t3_recovery.bin` created by the `-dump` process.

-scan

(Optional) Scan the specified MDisk or drive for system metadata.

-dump

(Required) Dump metadata from the specified MDisk or drive to file `/dumps/t3_recovery.bin`

-end end_arg

(Optional) The last LBA in which to look for metadata on the disk.

-file filename_arg

(Required) Specifies the file in which you want the results of a scan operation. The file is placed into the node in the directory `/dumps`, and can be retrieved using Secure Copy (**scp**). The file can be subsequently cleaned using the **cleardumps** command.

-disk UID_arg

(Required) Specifies the UID of the MDisk or drive that you want to scan, or remove a dump from.

-start start_arg

(Required) Specifies the following conditions:

- When used with **-scan**: The first LBA in which to look for metadata on the disk.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

Remember: You must specify one of the following:

- IPv4 system IP, gateway and subnet
- IPv6 system IP, gateway, and prefix

Description

This command creates a new system.

An invocation example using specific `-clusterip`, `-gw`, and `-mask` parameters

```
satask mkcluster -clusterip 192.168.1.2 -gw 192.168.1.1 -mask 255.255.255.0
```

The resulting output:

No feedback

mkcluster (Deprecated)

The **mkcluster** system command is deprecated. Use the **satask mkcluster** command to create a new clustered system (system).

satask overridequorum

Use the **overridequorum** command to invoke the manual override command.

Syntax

►►—satask overridequorum— —-force—►►

Parameters

-force

(Required) Overrides any quorum decision made by the system.

Important: Using this option might result in a loss of access. If used incorrectly this results in different nodes in the system using different copies of mirrored volumes simultaneously. Only use this command for disaster recovery scenarios where all nodes at one site have been lost.

Description

This command invokes the manual override command. This command is valid on nodes that are in starting state with either of the following node errors:

- 551
- 921

Remember: This command is only applicable when a system is configured as a HyperSwap or stretched system by specifying:

```
chsystem -topology
```

An invocation example

```
satask overridequorum
```

The resulting output:

No feedback

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.

Note: If `panel_name` is not supplied, this applies to the node on which the command is running.

-force

(required) 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 system code (code).

Description

This command starts automatic recovery for a specific node. Use this command only when the node reports corrupted code.

An invocation example

```
satask rescuenode -force 112233
```

The resulting output:

No feedback

resetpassword

Use the **resetpassword** command to reset the clustered system (system) superuser password to *passwOrd*.

Syntax

```
►► — satask — resetpassword — —————►◄
```

Parameters

satask

System Administrator task; service commands that are only used in specific circumstances.

Description

This command resets the system superuser password to *passwOrd*. You are prompted for a new password when you next log in to the graphical user interface (GUI).

An invocation example

`satask resetpassword`

The resulting output:

No feedback

restartservice

Use the **restartservice** command to restart a named service.

Syntax

```
▶▶ satask — restartservice — -service — service_name — panel_name —▶▶
```

Parameters

satask

Specifies a system administrator's task (such as service commands that are only used in specific circumstances).

-service *service_name*

(Required) Specifies the name of the service that you want to restart. The following services are supported:

sshd

Specifies a Secure Shell (SSH) Daemon.

slpd

Specifies a Service Location Protocol Daemon.

easy

Specifies Easy Tier service.

tomcat

Specifies a Web server.

cimom

Specifies CIMOM.

panel_name

(Optional) Identifies the node being serviced. If *panel_name* is not supplied, the command applies to the node on which the command is running.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

Description

Important: When directed to do so by your product support team, use this command to restart a named service.

An invocation example

`satask restartservice -service cimom`

The resulting output:

No feedback

(satask) setlocale

Use the **setlocale** command to change the **satask** and **sainfo** command output to the chosen language on the current node.

Syntax

►— satask — — 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 **satask setlocale** command to change the locale setting for the system; all interface output is changed to the chosen language. For example, to change the language to Japanese, type the following:

```
satask 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
- 8 Korean
- 9 Portuguese (Brazilian)
- 10 Russian

Note: This command does not change the front panel display panel settings.

An invocation example (where 3 is Japanese)

```
satask setlocale -locale 3
```

The resulting output:

No feedback

An invocation example (where 8 is Korean)

```
satask setlocale -locale 8
```

The resulting output:

No feedback

setpacedccu

Use the **setpacedccu** command to flag a node to participate in a user-paced system update.

Syntax

```
▶▶ — satask — setpacedccu — — panel_name —————▶▶
```

Parameters

panel_name

(Optional) Identifies the node being serviced.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

Description

Use this command to flag a node to participate in a user-paced system update. This command can only be used when the node is:

- In a service state
- Error-free
- Not part of a clustered system when the node is out of a service state

An invocation example

```
satask setpacedccu
```

The resulting output:

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

(Required) 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.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

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 use this command when completing 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.

An invocation example

```
satask settempsshkey
```

The resulting output:

No feedback

snap

Use the **snap** command to create a snap file on the node or enclosure that you specify.

Syntax

```
➤ satask — snap — [ -dump ] — [ -noimm ] — [ -noimm ] — [ panel_name ] ➤
```

Parameters

-dump

(Optional) Collects the existing dump.

-noimm

(Optional) Disables Integrated Management Module (IMM) First Failure Data Capture (FFDC) initiation and collection.

panel_name

(Optional) Identifies the node or enclosure being serviced.

Note: If panel_name is not supplied, this applies to the node or enclosure on which the command is running.

Description

This command creates a snap file on the node that you specify.

An invocation example

```
satask snap
```

The resulting output:

No feedback

Important: The name of the output file (placed on the specified node) is `snap.single.nodeid.date.time.tgz`.

An invocation example

```
satask snap -noimm
```

The resulting output:

No feedback

startservice

Use the **startservice** command to enter a service state.

Syntax

```
▶▶ satask — startservice — [ -force ] [ panel_name ] ▶▶
```

Parameters

satask

System Administrator task; service commands that are only used in specific circumstances.

-force

(Optional) Overrides checking of clustered system (system) membership.

Important: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information .

panel_name

(Optional) Identifies the node being serviced.

Note: If panel_name is not supplied, this applies to the node on which the command is running.

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 system 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.

An invocation example

```
satask startservice
```

The resulting output:

No feedback

stopnode

Use the **stopnode** command to power off, reboot, or warmstart a node.

Syntax

```
▶▶ satask — stopnode — [ -poweroff  
[ -reboot ]  
[ -warmstart ] ] [ panel_name ] ▶▶
```

Parameters

-poweroff

(Required if **-reboot** and **-warmstart** are not specified) Powers off the node.

-reboot

(Required if **-poweroff** and **-warmstart** are not specified) Reboots the node.

-warmstart

(Required if **-poweroff** and **-reboot** are not specified) Restarts the I/O process and issues a diagnostic dump (also known as a full dump).

Important: You can also specify `stopssystem -node -reset` to restart the I/O process (but in a more controlled manner).

panel_name

(Optional) Identifies the node being serviced.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

Description

Use the **stopnode** command to power off, reboot, or warmstart a node.

An example for powering off canister 1 in enclosure 2

```
satask stopnode -poweroff 02-1
```

The resulting output:

No feedback

An example for rebooting a node

```
satask stopnode -reboot
```

The resulting output:

No feedback

stopservice

Use the **stopservice** command to exit a service state.

Syntax

```

▶▶— satask — stopservice — ┬──────────────────────────────────────────────────────────────────────────────────▶▶
                             │
                             └─ panel_name ─┘

```

Parameters

panel_name

(Optional) Identifies the node being serviced.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

Description

This command exits service state entered using the **startservice** command, and exits the service state on the local node.

An invocation example

```
satask stopservice
```

The resulting output:

Node is no longer in service state

t3recovery

Use the **t3recovery** command to prepare and start a T3 recovery.

Syntax

```
►— satask — t3recovery — [ -prepare ] [ -execute ] [ panel_name ] ►
```

Parameters

panel_name

(Optional) Identifies the node being serviced.

Note: If *panel_name* is not supplied, this applies to the node on which the command is running.

-prepare

(Required) Search for T3 recovery data. This locates the date of the necessary backup file and quorum disk.

-execute

(Required) Start T3 recovery using recovered data.

Description

This command prepares and starts a T3 recovery.

Important: Progress of a T3 recovery can be displayed using the **sainfo lscmdstatus** command.

An invocation example

```
satask t3recovery -prepare
```

The resulting output:

No feedback

An invocation example

```
satask t3recovery -execute
```

The resulting output:

No feedback

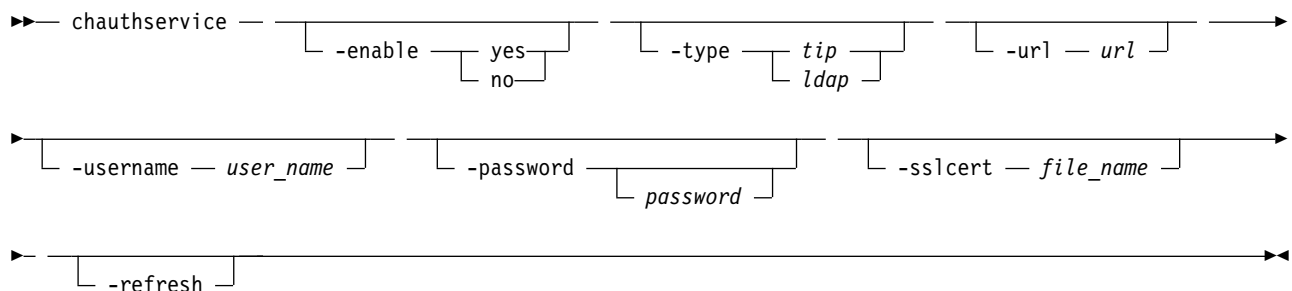
Chapter 28. User management commands

You can use the command-line interface (CLI) to configure remote authentication service and manage users and user groups on the clustered system.

chauthservice

Use the **chauthservice** command to configure the remote authentication service of the clustered system (system).

Syntax



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.

-type tip | ldap

(Optional) Specifies the authentication service type (TIP or native LDAP). An LDAP server must be configured. Before changing **-type**, ensure that the remote authentication type selected is properly configured.

Remember:

- The remote authentication service must be enabled (**-enable yes**) for this setting to come into effect.
- Before changing **-type** from *ldap* to *tip*, ensure that all users configured for remote authentication have both an SSH key and password configured.

-url url

(Optional - IBM Security Services only) Specifies the website address (URL) of Security Services, which is referred to as TIP in the CLI. The host part of the URL must be a valid numeric IPv4 or IPv6 network address. You can use the following characters in the URL:

- a - z
- A - Z
- 0 - 9
- -
- ~
- :
- [

-]
- %
- /

The maximum length of the URL 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. The certificate file must be in valid PEM format and have a maximum length of 12 kilobytes.

-refresh

(Optional) Causes the SAN Volume Controller to invalidate any remote user authorizations that are cached on the system. Use this when you modify user groups on the authentication service and want the change to immediately take effect on the SAN Volume Controller.

Note: If you clear the cache, anyone using the system might have to log in again (for example, if credentials are provided to one of the defined LDAP servers).

Description

This command can be used to select and enable a remote authentication service for use with the system.

The system can be configured to authenticate users against IBM Security Services or using Lightweight Directory Access Protocol (LDAP).

Before enabling remote authentication, ensure that the properties of the service are properly configured on the system. It is not necessary to disable the remote authentication service to change its properties. This command can be used to configure the TIP properties. LDAP authentication can be configured using the **chldap** command, and LDAP servers can be added to the system using the **mkldapserver** command.

Remember: For the authentication type to be set to LDAP with authorization enabled (true), an LDAP server must be configured. For authentication type to be set to TIP with authorization enabled (true), the TIP settings (URL, user name, password) must be configured.

When the authentication service is enabled or the configuration is changed, the system does not test whether the remote authentication system is operating correctly.

- To establish whether the system is operating correctly, issue the **lscurrentuser** command for a remotely authenticated user. If the output lists the user 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.

- To establish whether LDAP is operating correctly, in addition to the **lscurrentuser** command, issue the **testldapserver** command. The **testldapserver** command can be issued whether or not remote authentication is enabled, and can be used to test the connection to LDAP servers, as well as user authorization and authentication.

The website address in the TIP *url* parameter 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
```

To disable the remote authentication service in a controlled manner when it is not available, use the **enable** parameter with the **no** option.

An invocation example

To fully configure and enable authentication with IBM Security Services:

```
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:

```
chauthservice -enable no
```

The resulting output:

No feedback

An invocation example

To refresh the SAN Volume Controller remote authorization cache:

```
chauthservice -refresh
```


The resulting output:

No feedback

chcurrentuser

Use the **chcurrentuser** command to change the attributes of the current user.

Syntax



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

```
chcurrentuser -password secret -nokey
```

The resulting output:

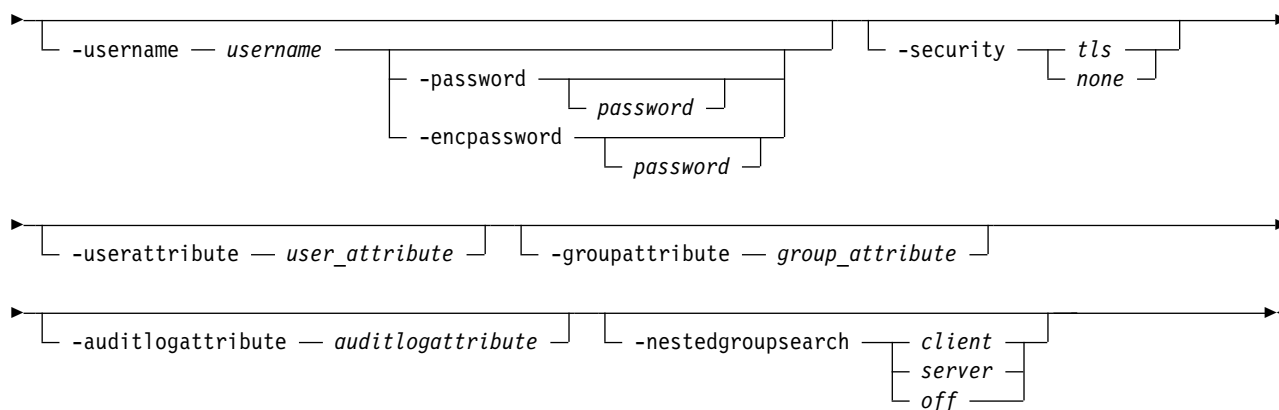
No feedback

chldap

Use the **chldap** command to change system-wide Lightweight Directory Access Protocol (LDAP) configuration. This command can be used to configure remote authentication with LDAP. These settings apply when authenticating against any of the LDAP servers configured using the **mkldapserver** command.

Syntax





Parameters

-type *ad | itds | other* | **-reset**

(Optional) Specify the LDAP server type, or reset LDAP configuration to defaults for the current server type. Defaults for the configured server type:

- Active Directory (AD)
- IBM Tivoli Directory Server (ITDS)
- Other

-username *username*

(Optional) Specifies a username for administrative binding. This can be:

Note:

- A distinguished name (DN)
- A user principal name (UPN) or NT login name for Active Directory

-password *password*

(Optional) Specifies the password for the administrative binding. You can optionally specify the password with this 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.

-encpassword *password*

(Optional) Specifies the password for the enclosure. You can optionally specify the password with this 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.

-security *tls | none*

(Optional) Specifies the type of security to use when communicating with LDAP servers.

-userattribute *user_attribute*

(Optional) Specifies the LDAP attribute used to determine the user name of remote users. The user attribute must exist in your LDAP schema and must be unique for each of your users.

-groupattribute *group_attribute*

(Optional) Specifies the LDAP attribute used to determine the group memberships of remote users. The attribute must contain either the DN of a group or a colon-separated list of group names.

-auditlogattribute *auditlogattribute*

(Optional) Specifies the LDAP attribute used to determine the identity of remote users. When a user performs an audited action, this information is recorded in the audit.

-authcacheminutes *auth_cache_minutes*

(Optional) Specifies the period for which to cache authentication details.

-nestedgroupsearch *client* | *server* | *off*

(Optional) Specifies whether nested groups are evaluated on the client (clustered system), server (authentication service), or are not evaluated not at all.

Description

At least one parameter must be specified.

The **chldap** command can be run whether or not LDAP authentication is enabled. Specifying **-reset** or **-type** populates the default values unless otherwise specified.

You can only specify **-password** or **-encpassword** if **-username** is specified.

The **-type** parameter values are only set to defaults for the specified type if the type is different from the existing type.

If the type is **itds**, **-nestedgroupsearch** cannot be executed (nested groups are evaluated by default). If the type is **ad**, **-nestedgroupsearch** can only be set to **client** or **off** because there is no server support. If the type is **other**, the **-nestedgroupsearch** parameter is fully configurable.

Use **-username** to specify a distinguished name (DN), user principal name (UPN), or NT login name. Distinguished names (DN) must be a sequence of attribute=value pairs separated by a comma (,), semi-colon(;), or plus sign (+). A backslash (\,) must be used to escape special characters, and can also be used to specify UTF-8 characters using their byte encoding. For example, c acute can be represented as \C4\87. NT logins are valid for only the Active Directory and must be in the DOMAIN\user format. These logins must not start or end with a period (.) and both the DOMAIN and the user must not use the following characters: \/:?<>| UPN logins are valid for Active Directory only and must be in the format user@suffix. Both user and suffix not use spaces or the following characters: ()<>,;:\>[]@

Tip:

- Remember that **-userattribute**, **-groupattribute**, and **-auditlogattribute** accept values that:
 - Must begin with a letter
 - Only contain ASCII letters, digit characters, and hyphens
 - Are case-insensitive

The following LDAP (first-time) configuration suggestions assist with LDAP server setup:

Important:

- Ensure that the system is configured appropriately according to your LDAP schema. Issue **chldap -type** to populate the system's LDAP configuration with the server type defaults. Issue **chldap -reset** to return to these defaults at any time.
 - (Advanced) For all server types, users are authenticated with a username configured in the LDAP attribute **user_attribute**. This attribute must exist in the LDAP schema and must be unique for each user. It is configurable by issuing **chldap -userattribute**. Active Directory users can also authenticate using their UPN or NT login names.
 - (Advanced) Authenticated users are assigned roles according to their LDAP group memberships. Each user's group memberships must be stored in the LDAP attribute **group_attribute**. This can be either an LDAP attribute containing the DN of the user's LDAP group, or an LDAP attribute containing a colon-separated list of user group names. It is configurable by issuing **chldap -groupattribute**.
 - (Advanced) When an LDAP authenticated user runs a command that is audited, the user's login name is placed in the audit log. The name is extracted from the LDAP attribute **audit_log_attribute**, which is configurable by issuing **chldap -auditlogattribute**.

- Ensure that the system is able to search within the user and group trees on LDAP servers. By default the system authenticates anonymously. Consequently, you must either permit anonymous searches of the LDAP directory, or create an LDAP user with the appropriate permissions and issue the **chldap -username** and **chldap -password** commands to instruct the system to search as this user.
- Ensure that the system is able to connect with the appropriate level of security. Passwords are sent to the LDAP server as clear text, so Transport Layer Security (TLS) encryption is recommended. Issue **chldap -security** to change the security level.
- (Advanced): On Active Directory and some other LDAP servers, the system (by default) identifies groups to which users belong directly. To assign users permissions according to a parent group, enable the nested group search on the client by issuing **chldap -nestedgroupsearch**. This setting has an additional performance overhead and supports up to 8 levels of nesting.

An invocation example

```
chldap -type
itds -username uid=joebloggs,cn=admins,dc=company,dc=com -password passw0rd
-auditlogattribute descriptiveName
```

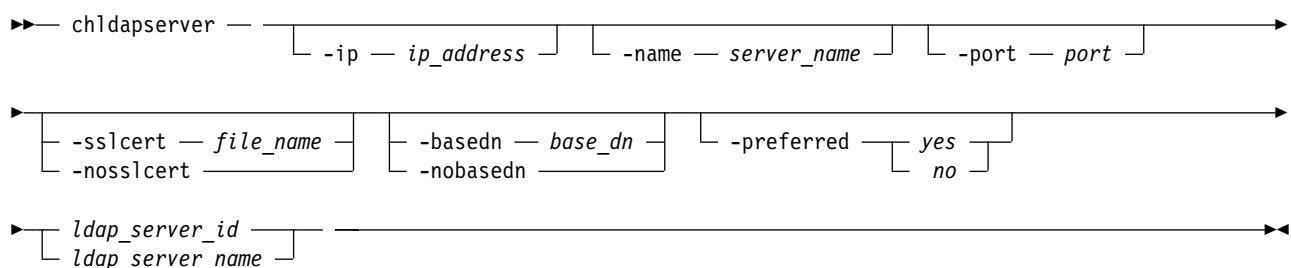
The resulting output:

No feedback

chldapserver

Use the **chldapserver** command to modify a Lightweight Directory Access Protocol (LDAP) server.

Syntax



Parameters

- ip ip_address**
(Optional) Specifies the server IP address (Internet Protocol Version 4 or 6).
- name server_name**
(Optional) Specifies the LDAP server name.
- port port**
(Optional) Specifies the LDAP server port.
- sslcert file_name | -nossllcert**
(Optional) Set (**-sslcert**) or clear (**-nossllcert**) the secure socket layer (SSL) certificate.
- basedn base_dn | -nobasedn**
(Optional) Use the base distinguished name (DN) for search (**-nobasedn** indicates to use the default DN).
- preferred yes | no**
(Optional) Specifies whether the server is preferred over other configured LDAP servers (or not preferred).

ldap_server_id | ldap_server_name
(Required) Specifies the LDAP server ID or name.

Description

Important: During normal operation, LDAP requests are sent to **-preferred** servers depending on availability. If no servers are marked as **-preferred**, LDAP requests are sent to configured servers based on availability.

If **-sslcert** is specified, the server certificate is verified while authenticating. The SSL certificate must exist on the current node. If **-nossllcert** is specified, any certificate file is deleted and the server certificate is not checked.

The **-basedn** parameter indicates the distinguished name (DN) to use as a base from which to search for users in the LDAP directory. If Transport Layer Security (TLS) is enabled and **-sslcert** is specified, the server certificate is verified during authentication. The secure socket layer (SSL) certificate must exist on the node being used. Otherwise, a server certificate is not checked.

The clustered system (system) must be configured with an appropriate version IP address when **-ip** is specified. The IP address specified with the **-ip** parameter must be of a version supported by the system. The certificate file must be in valid PEM format and have a maximum length of 12 kilobytes.

Distinguished names must be a sequence of attribute=value pairs separated by a comma (,), semi-colon(;), or plus sign (+) escaping special characters with \ where appropriate, and specified UTF-8 characters using their byte encoding. For example, , for commas or \C4\87 for the UTF-8 character c acute.

This command runs whether or not LDAP authentication is enabled.

Remember: There can be a maximum of six configured LDAP servers. If you attempt to create a seventh LDAP server an error is returned.

An invocation example with basic server details

```
chldapserver -ip 192.135.60.3 -port 400 ldapserver0
```

The resulting output:

No feedback

An invocation example specifying an SSL certificate

```
chldapserver -sslcert /tmp/activedirectorycert.pem 0
```

The resulting output:

No feedback

An invocation example to remove an SSL certificate

```
chldapserver -nossllcert 0
```


The resulting output:

No feedback

chuser

Use the **chuser** command to change 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 **yes** or **no** must be set.

-usergrp *group_name* | *group_id*

(Optional) Specifies the new group for the user.

user_name | *user_id*

(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

```
chuser -remote no -usergrp Monitor -nokey jane
```

The resulting output:

No feedback

chusergrp

Use the **chusergrp** command to change the attributes of an existing user group.

Syntax

```
➤ chusergrp — [-remote yes|no] [group_id|group_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.

group_id | *group_name*

(Required) The ID or name of 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

```
chusergrp -role Administrator admin
```

The resulting output:

No feedback

lscurrentuser

Use the **lscurrentuser** command to display the name and role of the logged-in user.

Syntax

```
➤ 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.

lsldap

Use the **lsldap** command to display the details for the system-wide Lightweight Directory Access Protocol (LDAP) configuration.

Syntax

```
►► lsldap — [ -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, with the width of each column set to the maximum possible width of each item of data. In a detailed view, each item of data is an individual row, and if displaying headers, 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. Enter **-delim :** on the command line, and 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 specified *delimiter* separates the data from its header

Description

Table 101 provides the attribute values that can be displayed as output view data.

Table 101. *lsldap attribute values*

Attribute	Value
type	LDAP server type: <ul style="list-style-type: none">• Active Directory: ad• IBM Tivoli Directory Server: itds• Other: other
enabled	Is native LDAP authentication enabled?
error_sequence_number	Sequence number of non-fixed LDAP configuration error log
username	Binding username or distinguished name (or blank if there is none)
security	Type of security in use: <ul style="list-style-type: none">• Transport Layer Security: tls• No security: none
user_attribute	LDAP attribute representing user login
group_attribute	LDAP attribute representing user group membership
audit_log_attribute	LDAP attribute representing user name in audit log
auth_cache_minutes	Period (in minutes) for which to cache session details
nested_group_search	Handling of nested groups: <ul style="list-style-type: none">• No nested group handling: off• Search nested groups on the client: client• Search nested groups on the server: server

An invocation example

```
lsldap -delim :
```

The resulting output:

```
type:ad
enabled:yes
error_sequence_number:12
username:admin@company.com
security:tls
user_attribute:sAMAccountName
group_attribute:memberOf
audit_log_attribute:userPrincipalName
auth_cache_minutes:10
nested_group_search:off
```

lsldapserver

Use the **lsldapserver** command to display the most recent details for all configured Lightweight Directory Access Protocol (LDAP) servers.

Syntax

```
➤➤ lsldapserver — [ -nohdr ] [ -delim delimiter ] [ ldap_server_id ] [ ldap_server_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, with the width of each column set to the maximum possible width of each item of data. In a detailed view, each item of data is an individual row, and if displaying headers, 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. Enter **-delim :** on the command line, and 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 specified *delimiter* separates the data from its header

ldap_server_id | ldap_server_name

(Optional) Specifies the ID or name for LDAP server being used.

Description

Remember:

- The base distinguished name (DN) is located at the end of the concise view information; other fields must be added before the base DN.
- The command fails if a server is specified that does not exist.

Table 102 provides the attribute values that can be displayed as output view data.

Table 102. lslldapserver attribute values

Attribute	Value
id	ID of the LDAP server
name	Name of the LDAP server
error_sequence_number	Sequence number of non-fixed LDAP server error log
IP_address	IP address of the LDAP server (Internet Protocol Versions 4 and 6)
port	LDAP server port
cert_set	Certificate setting (Is a certificate configured?)
preferred	Server preference (Is this server preferred?)
base_dn	Base distinguished name used in LDAP searches

Description

This command displays details for the configured Lightweight Directory Access Protocol (LDAP) servers.

Note: There is a maximum of six configured LDAP servers.

A concise invocation example

```
lslldapserver -delim :
```

The resulting output:

```
id:name:error_sequence_number:IP_address:port:cert_set:preferred:base_dn
0:ldapsrv0::192.135.60.3:389:no:yes:ou=users,dc=company,dc=com
1:ldapsrv1:12:192.135.60.4:389:no:no:ou=users,dc=company,dc=com
2:ldapsrv2::192.135.60.5:389:yes:yes:ou=users,dc=company,dc=com
3:ldapsrv3::192.135.60.6:389:yes:no:ou=users,dc=company,dc=com
```

A detailed invocation example

```
lsldapsrv -delim : ldapsrv0
```


The resulting output:

```
id:0
name:ldapsrv0
error_sequence_number:
IP_address:192.135.60.3
port:389
cert_set:no
preferred:yes
base_dn:ou=users,dc=company,dc=com
```

lsuser

Use the **lsuser** command to display a list of the users that have been created on the clustered system (system).

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 (*).
- The command can contain a maximum of one wildcard.

- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:
`lsuser -filtervalue "usergrp_name=md*"`

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalueattribute=value** parameter:

- password
- ssh_key
- remote
- usergrp_id
- usergrp_name

`usergrp_name | usergrp_id`

(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 system.

A concise invocation example

```
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		

A detailed invocation example

```
lsuser 1
```

The resulting output:

```
id 1
name tpc_admin
password yes
ssh_key no
remote no
usergrp_id 0
usergrp_name SecurityAdmin
```

lsusergrp


Use the **lsusergrp** command to display a list of the user groups that have been created on the clustered system (system).

Syntax

```

▶▶ lsusergrp — [ -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.

-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:

```
lsusergrp -filtervalue "role=md*"
```

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalue attribute=value** parameter:

- role_id
- role_name
- remote

usergrp_name | usergrp_id

(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 system.

An invocation example

```
lsusergrp
```

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

mkldapserver

Use the **mkldapserver** command to display the data used to create a Lightweight Directory Access Protocol (LDAP) server.

Syntax

```

▶▶ mkldapserver — — -ip — ip_address — [ -name — server_name ] [ -port — port ]
▶ [ -sslcert — file_name ] [ -basedn — base_dn ] [ -preferred ]

```

Parameters

- ip** *ip_address*
(Required) Specifies the server IP address (Internet Protocol Version 4 or 6).
- name** *server_name*
(Optional) Specifies the LDAP server name.
- port** *port*
(Optional) Specifies the LDAP server port.
- sslcert** *file_name*
(Optional) Set the SSL certificate.
- basedn** *base_dn*
(Optional) Use the base distinguished name for search.
- preferred**
(Optional) Specifies that this server is preferred over other configured LDAP servers.

Description

Important: During normal operation, LDAP requests are sent to **-preferred** servers depending on availability. If no servers are marked as **-preferred**, LDAP requests are sent to configured servers based on availability.

If **-sslcert** is specified, the server certificate is verified while authenticating.

Note: The SSL certificate must exist on the current node.

The **-basedn** parameter indicates the distinguished name (DN) to use as a base from which to search for users in the LDAP directory. If Transport Layer Security (TLS) is enabled and **-sslcert** is specified, the server certificate is verified during authentication. The secure socket layer (SSL) certificate must exist on the node being used, otherwise a server certificate is not checked.

The clustered system (system) must be configured with an appropriate version IP address when **-ip** is specified. The IP address specified with the **-ip** parameter must be of a version supported by the system. The certificate file must be in valid PEM format and have a maximum length of 12 kilobytes.

Distinguished names must be a sequence of attribute=value pairs separated by a comma (,), semi-colon(;), or plus sign (+) escaping special characters with a backslash (\) where appropriate, and specified UTF-8 characters using their byte encoding. For example, \, for commas or \C4\87 for the UTF-8 character c acute.

This command runs whether or not LDAP authentication is enabled.

Remember: There is a maximum of six configured LDAP servers. Attempting to create a seventh LDAP server returns an error.

An invocation example

```
mkldapserver -ip 192.135.60.3
```

The resulting output:

```
LDAP Server, id [0], successfully created
```

mkuser

Use the **mkuser** command to create either a local or a remote user to access a SAN Volume Controller clustered system (system).

Syntax

```
mkuser -name user_name
  -remote
  -usergrp group_id group_name -keyfile sshkey_filename -password cleartext_password
```

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**

(Required) Specifies whether the user authenticates to the system using a remote authentication service or system authentication methods. Either the **remote** parameter or the **usergrp** parameter must be set. If **usergrp** is specified, it must be followed by *group_name* or *group_id* (see next parameter).

group_name | *group_id*

(Required if **usergrp** is specified) The ID or name of the user group with which the local user is to be associated.

-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 system. 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 system. 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.

Up to 400 users can be defined on the system. You can also create new users and assign keys to them.

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 system state and activated for the user, and the input file is deleted.

An invocation example

```
mkuser -name jane -usergrp Service -password secret
```

The resulting output:

```
User, id [1], successfully created
```

mkusergrp

Use the **mkusergrp** command to create a new user group.

Syntax

```
➤ mkusergrp — — -name — group_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: %: ",*' .

-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 clustered system by role. Use the **lsusergrp** command to view a list of user groups that have been created on the clustered system.

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

You can issue any information display command and, additionally, the following commands:

- **finderr**
- **dumperrlog**
- **dumpinternallog**
- **chcurrentuser**
- **ping**
- **svconfig backup**

CopyOperator

You can issue the following commands:

- **prestartfcconsistgrp**
- **startfcconsistgrp**
- **stopfcconsistgrp**
- **chfcconsistgrp**
- **prestartfcmap**
- **startfcmap**
- **stopfcmap**
- **chfcmap**
- **starttrcconsistgrp**
- **stoptrcconsistgrp**
- **switchtrcconsistgrp**
- **chrccconsistgrp**
- **starttrcrelationship**
- **stoptrcrelationship**
- **switchtrcrelationship**
- **chrcrelationship**
- **chpartnership**

In addition, you can issue all of the commands allowed by the **Monitor** role.

Service

You can issue the following commands:

- **applysoftware**
- **setlocale**
- **addnode**
- **rmnode**
- **cherrstate**
- **writesernum**
- **detectmdisk**
- **includemdisk**
- **clearerrlog**
- **cleardumps**
- **settimezone**
- **stopsystem**
- **startstats**

- **stopstats**
- **settime**

In addition, you can issue all of the commands allowed by the **Monitor** role.

Administrator

You can issue any command other than:

- **chauthservice**
- **mkuser**
- **rmuser**
- **chuser**
- **mkusergrp**
- **rmusergrp**
- **chusergrp**
- **setpwdreset**

VASAProvider

The system uses this software uses this role to implement the VMware Virtual Volumes function. (It provides a group with users that can be used by that software.) You can issue any command other than:

- **chauthservice**
- **chldap**
- **chldapserver**
- **chsecurity**
- **chuser**
- **chusergrp**
- **mkldapserver**
- **mkuser**
- **mkusergrp**
- **rmldapserver**
- **rmuser**
- **rmusergro**
- **setpwdreset**

SecurityAdmin

You can issue all commands.

The command returns the ID of the created user group.

An invocation example

```
mkusergrp -name support -role Service
```

The resulting output:

```
User Group, id [5], successfully created
```

An invocation example

```
mkusergrp -role VasaProvider -name myVasaProvider
```

The resulting output:

```
User Group, id [5], successfully created
```

rmldapserver

Use the **rmldapserver** command to delete a Lightweight Directory Access Protocol (LDAP) server.

Syntax

```
➤— rmldapserver — [ ldap_server_id | ldap_server_name ] —➤
```

Parameters

ldap_server_id | ldap_server_name
(Required) Specifies the LDAP server ID or name to delete.

Description

Remember:

- If remote authentication with LDAP is enabled, the final LDAP server cannot be deleted. To delete the final LDAP server disable LDAP authentication by specifying **chauthservice -enable no**.
- The **rmldapserver** command can be specified whether or not LDAP authentication is enabled.

An invocation example

```
rmldapserver ldapserver0
```

The resulting output:

No feedback

rmuser

Use the **rmuser** command to delete a user.

Syntax

```
➤— rmuser — [ user_id | user_name ] —➤
```

Parameters

user_id or user_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 modify a user.

An invocation example

```
rmuser jane
```

The resulting output:

No feedback

rmusergrp

Use the **rmusergrp** command to delete a user group.

Syntax

```
►►—rmusergrp— [—force—] [group_id | group_name]—►►
```

Parameters

-force

(Optional) Specifies that the user group should be deleted even if there are users in the group.

Important: Using the force parameter might result in a loss of access. Use it only under the direction of the your product support information.

group_id | group_name

(Required) 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

```
rmusergrp support
```

The resulting output:

No feedback

testldapserver

Use the **testldapserver** command to test a Lightweight Directory Access Protocol (LDAP) server.

Syntax

```
►►— testldapserver — — -delim — delimiter — [—username — user_name — [—password — password —]] —►►
```



```
► [ldap_server_id | ldap_server_name] —►
```

Parameters

-delim delimiter

(Optional) By default, in a concise view all columns of data are space-separated, with the width of

each column set to the maximum possible width of each item of data. In a detailed view, each item of data is an individual row, and if displaying headers, 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. Enter **-delim :** on the command line, and 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 specified *delimiter* separates the data from its header

-username *user_name*

(Optional) Specifies the user name to test.

-password *password*

(Optional) Specifies the password to test. You can optionally specify the password with this 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.

Note: The **-password** parameter is only valid if **-username** is specified. The actual password does not need to be supplied.

ldap_server_id | *ldap_server_name*

(Optional) Specifies the LDAP server ID or name to test.

Description

The **testldapserver** command allows three levels of testing:

- Server connection test (issue **testldapserver** without supplying username or password). This verifies that a connection can be established with the server while authenticating using the configured administrator credentials according to the LDAP configuration.
- Server connection, LDAP configuration, and user authorization test (issue **testldapserver** with a username). This verifies that:
 - A connection can be established with the server while authenticating using the configured administrator credentials.
 - The LDAP attributes are correctly configured on the system.
 - The user has been assigned a role.
- Server connection, LDAP configuration, and user authentication test (issue **testldapserver** with a username and password). This verifies that:
 - A connection can be established with the server while authenticating using the configured administrator credentials.
 - The user authenticates with the supplied password

No specific server errors indicates success.

Important: This command works whether or not LDAP authentication is selected or enabled with the **chauthservice** command.

Table 103 provides the attribute values that can be displayed as output view data.

Table 103. testldapserver attribute values

Attribute	Value
id	LDAP server ID
name	LDAP server name
error	Critical server error (or success, depending on situation) encountered

An invocation example with one LDAP server and no specific user information

```
testldapserver -delim ":" ldapserver1
```

The resulting output:

```
id:name:error
1:ldapserver1:CMMVC7075I The LDAP task completed successfully
```

An invocation example with all LDAP servers using a UPN

```
testldapserver -username bloggs@company.com -delim ":"
```

The resulting output:

```
id:name:error
0:ldapserver0:CMMVC6518E The task has failed because no roles
                        are defined for the current user on the system.
1:ldapserver1:CMMVC7075I The LDAP task completed successfully.
2:ldapserver2:CMMVC7075I The LDAP task completed successfully.
```

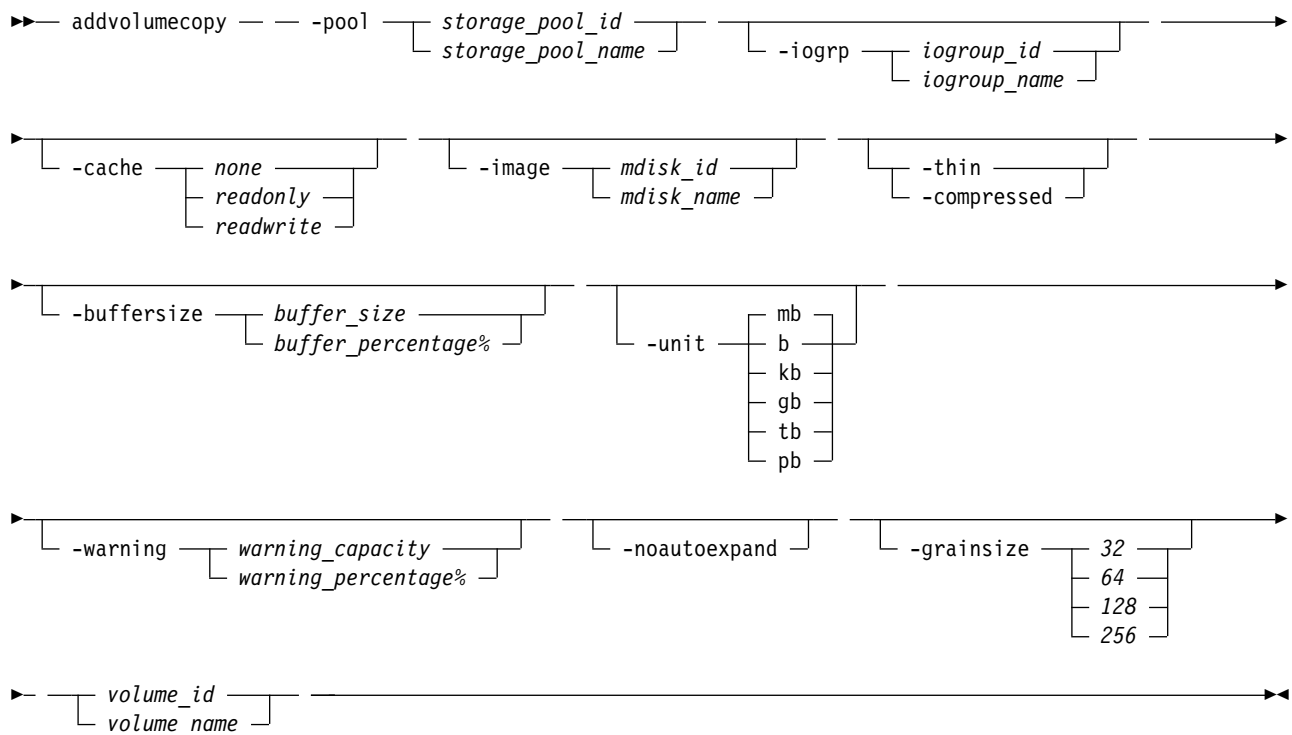
Chapter 29. Volume commands

The volume commands enable you to work with volume options with the SAN Volume Controller.

addvolumecopy

Use the **addvolumecopy** command to add a copy to an existing volume. On a standard topology system, you can use this command to add a mirrored copy to an existing volume. On a stretched or HyperSwap topology system, you can use this command to convert an existing basic volume into a highly available volume by adding a copy of that volume at a second site.

Syntax



Parameters

-pool *storage_pool_id* | *storage_pool_name*
(Required) Specifies the storage pool in which to create the new volume copy.

Remember: For stretched and hyperswap topology systems, the site of storage pool must not be the same as the site of the existing volume copy.

-iogrp *iogroup_id* | *iogroup_name*
(Optional) Specifies the I/O group that the new volume copy is cached in.

Note: This parameter applies only when you are creating a HyperSwap volume and requires that the system topology be hyperswap.

The I/O group must be in the same site as the storage pool that is being specified.

-cache *none* | *readonly* | *readwrite*

(Optional) Specifies the caching options for the volume copy. Valid entries are:

- *readwrite* enables the cache for the volume *
- *readonly* disables write caching but allows read caching for a volume
- *none* disables the cache mode for the volume

Note: This parameter applies only when you are creating a HyperSwap volume and requires that the system topology be hyperswap.

-image *mdisk_id* | *mdisk_name*

(Optional) Specifies that the volume copy is to be created (on any topology) in image mode and it specifies which currently unused MDisk to use.

Note: For a stretched or hyperswap topology system the MDisk site must match the storage pool site. If the storage pool is empty, the MDisk site must be 1 or 2, and the MDisk site cannot be the same as the site of the existing volume copy.

-thin

(Optional) Specifies that the volume copy is to be created with thin-provisioning. You cannot specify this parameter with **-compressed**.

Note: If you do not specify either **-thin** or **-compressed** the volume copy is created as one that is fully-allocated.

-compressed

(Optional) Specifies that the volume copy is to be created compressed. You cannot specify this parameter with **-thin**.

-buffersize *buffer_size* | *buffer_percentage%*

(Optional) Specifies the pool capacity the volume attempts to reserve as a buffer for thin-provisioned and compressed volumes. You must specify either **-thin** or **-compressed** with this parameter.

-warning *warning_capacity* | *warning_percentage%*

(Optional) Specifies a threshold at which a warning error log is generated for the volume copy. A warning is generated when the used disk capacity on the thin-provisioned or compressed copy exceeds the specified threshold. You can specify the threshold by using *warning_capacity* to specify a size, which defaults to MB unless the **-unit** parameter is specified.

Note: You can also specify a *warning_percentage%*, to use a percentage of the volume size. If you do not specify a warning threshold, a default value of 80% is used. To disable warnings, specify 0.

You must specify either **-thin** or **-compressed** with this parameter.

-unit *b* | *kb* | *mb* | *gb* | *tb* | *pb*

(Optional) Specifies the data units for the **-buffersize** and **-warning** parameters.

-noautoexpand

(Optional) Specifies that the volume copy not automatically expand as it is written to; the available buffer capacity decreases as the used capacity increases. The copy goes offline if the buffer capacity is consumed.

The buffer capacity can be increased by specifying `expandvdiskspace -rsize`. You must specify either **-thin** or **-compressed** with this parameter. If you do not specify this keyword the copy automatically expands as it is written to.

-grainsize *32* | *64* | *128* | *256*

(Optional) Sets the grain size (KB) for a thin-provisioned volume. If you are using the thin-provisioned volume in a FlashCopy map, use the same grain size as the map grain size for best performance. If you are using the thin-provisioned volume directly with a host system, use a small grain size. The grain size value must be 32, 64, 128, or 256 KB. The default is 256 KB.

volume_id | *volume_name*
(Required) Specifies the volume to add the volume copy to.

Description

Use the **addvolumecopy** command to add a copy to an existing volume. The new volume copy is synchronized with the current copy.

Note: A volume cannot have volume copies in different storage pools if cloud backup is enabled on the volume.

On a standard topology system, you can use this command to add a mirrored copy to an existing volume. On a stretched or HyperSwap topology system, you can use this command to convert an existing basic volume into a highly available volume by adding a copy of that volume to a second site.

A volume copy cannot be created in the same site as an existing copy of the volume. This command automatically adds the caching I/O group to the access I/O group set of the volume.

Add a volume copy to an existing volume

```
addvolumecopy -pool 2 volume5
```

The detailed resulting output:

No feedback

Add a thin volume copy to an existing volume

```
addvolumecopy -pool site2pool1 -thin 0
```

The detailed resulting output:

No feedback

Add a fully-allocated image-mode volume copy

```
addvolumecopy -image mdisk12 -pool 3 volume2
```

The detailed resulting output:

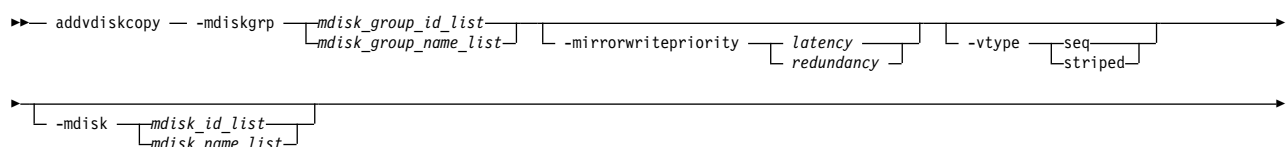
No feedback

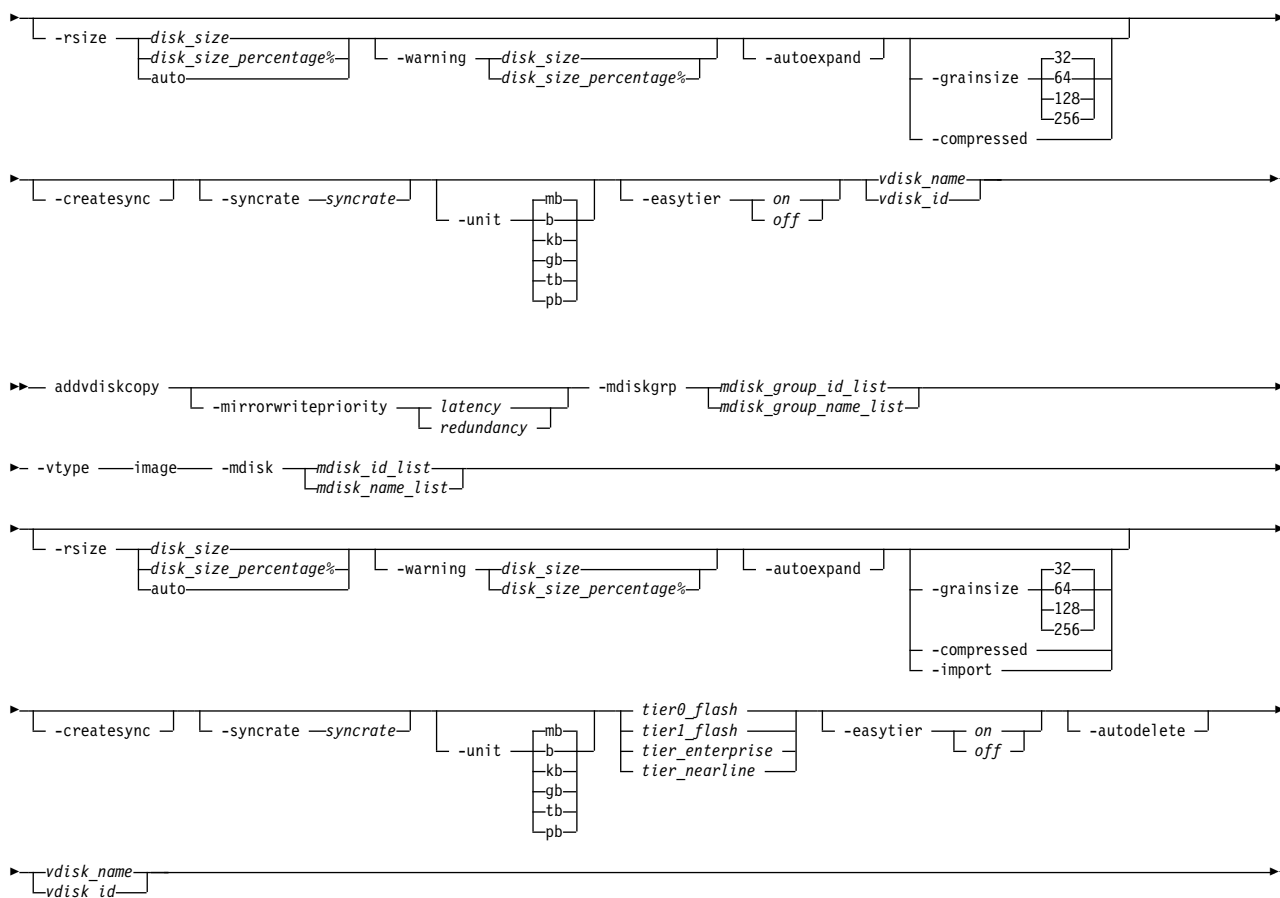
addvdiskcopy

Use the **addvdiskcopy** command to add a copy to an existing volume, which changes a nonmirrored volume into a mirrored volume. On a system with a hyperswap topology, use the **addvolumecopy** command to convert an existing volume to a hyperswap volume by adding a copy at a second site.

Note: The first syntax diagram depicts the addition of a sequential or striped mode volume. The second syntax diagram depicts the addition of an image mode volume.

Syntax





Parameters

-mdiskgrp *mdisk_group_id_list* | *mdisk_group_name_list*

(Required) Specifies the storage pools to use to create copies for the volume. You must specify a group for each copy that is being added.

Note: If the MDisk group is from a child pool, **-vtype** must be striped.

-mirrorwritepriority *latency* | *redundancy*

(Optional) Specifies how to configure the mirror write algorithm priority.

1. Choosing *latency* means a copy that is slow to respond to a write input/output (I/O) becomes unsynchronized, and the write I/O completes if the other copy successfully writes the data.
2. Choosing *redundancy* means a copy that is slow to respond to a write I/O synchronizes completion of the write I/O with the completion of the slower I/O to maintain synchronization.

-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 volume. The default virtualization type is **striped**.

Note: You cannot create an image or sequential mode volume from a child pool.

-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 *syncrate*

(Optional) Specifies the copy synchronization rate. A value of zero (0) prevents synchronization. For the supported **-syncrate** values and their corresponding rates, see Table 105 on page 660.

If not specified, the current value is unchanged.

-createsync

(Optional) Suppresses the synchronization of the new volume 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.

Note: You cannot specify **-createsync** for a volume that is fast formatting.

-rsize *disk_size* | *disk_size_percentage%* | **auto**

(Optional) Makes the copy thin-provisioned 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 volume copy that uses the entire size of the MDisk; if you specify the **-rsize auto** option, you must also specify the **-vtype image** option.

-compressed

(Optional) Adds exactly one copy to an existing volume that already has (only) one copy a volume, and enables compression. Requires the **-rsize** parameter also be specified.

Remember:

- You cannot specify this parameter with the **-grainsize** parameter.
- When you specify this parameter with the **-import** parameter, you must specify **-rsize auto**.

-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 thin-provisioned 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 volume size. If **-autoexpand** is enabled, the default value for **-warning** is 80% of the volume 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 thin-provisioned copies automatically expand their real capacities by allocating new extents from their storage pool. 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 storage pool 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 thin-provisioned volume copy. The grain size value must be 32, 64, 128, or 256 KB. The default is 256 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 thin-provisioned volume into the clustered system (system). Requires that the **-rsize** and **-vtype image** parameters also be specified.

-tier *tier0_flash* | *tier1_flash* | *tier_enterprise* | *tier_nearline*

(Optional) Specifies the MDisk tier when an image mode copy is added.

tier0_flash

Specifies a tier0_flash hard disk drive or an external MDisk for the newly discovered or external volume.

tier1_flash

Specifies an tier1_flash (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

tier_enterprise

Specifies a tier_enterprise hard disk drive or an external MDisk for the newly discovered or external volume.

tier_nearline

Specifies a tier_nearline hard disk drive or an external MDisk for the newly discovered or external volume.

-easytieron | off

(Optional) Determines if the IBM Easy Tier function is allowed to move extents for this volume. If a volume copy is striped and not being migrated the following table applies:

Table 104. 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 659)
Off	One	On	inactive (see note 2 on page 659)
Off	Two	Off	inactive (see note 2 on page 659)
Off	Two	On	inactive (see note 2 on page 659)
Measure	One	Off	measured (see note 3 on page 659)
Measure	One	On	measured (see note 3 on page 659)
Measure	Two	Off	measured (see note 3 on page 659)
Measure	Two	On	measured (see note 3 on page 659)
Auto (see note 6 on page 659)	One	Off	measured (see note 3 on page 659)
Auto (see note 6 on page 659)	One	On	measured (see note 3 on page 659)
Auto (see note 6 on page 659)	Two	Off	balanced (see note 4 on page 659)
Auto (see note 6 on page 659)	Two	On	active (see note 1 on page 659)
On (see note 6 on page 659)	One	Off	measured (see note 3 on page 659)
On (see note 6 on page 659)	One	On	balanced (see note 4 on page 659)
On (see note 6 on page 659)	Two	Off	measured (see note 3 on page 659)

Table 104. 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
On (see note 6)	Two	On	active (see note 1)
<p>Note:</p> <ol style="list-style-type: none"> 1. If the volume copy is in image or sequential mode or is being migrated, then the volume copy Easy Tier(tm) status is measured instead of active. 2. When the volume copy status is inactive, no Easy Tier(tm) functions are enabled for that volume copy. 3. When the volume copy status is measured, the Easy Tier(tm) function collects usage statistics for the volume but automatic data placement is not active. 4. When the volume copy status is balanced, the Easy Tier(tm) function enables performance-based pool balancing for that volume copy. 5. When the volume copy status is active, the Easy Tier(tm) function operates in automatic data placement mode for that volume. 6. The default Easy Tier(tm) setting for a storage pool is auto, and the default Easy Tier(tm) setting for a volume copy is on. This means that Easy Tier(tm) functions except pool performance balancing are disabled for storage pools with a single tier, and that automatic data placement mode is enabled for all striped volume copies in a storage pool with two or more tiers. 			

-autodelete

(Optional) Specifies the primary copy is deleted after the secondary copy is synchronized.

jvdisk_name | **vdisk_id**

(Required) Specifies the volume to add the volume copy to, either by ID or by name.

Description

The **addvdiskcopy** command adds a copy to an existing volume, which changes a nonmirrored volume into a mirrored volume. Use the **mkdiskgrp** parameter to specify the storage pools that provide storage for the copy; the **lsmdiskgrp** command lists the available storage pools and the amount of available storage in each group.

The **addvdiskcopy** command can be specified with a file system volume, but must be used with the same storage pool for that volume.

Remember: Only compressed copies are allowed to be added to file system volumes.

The **addvdiskcopy** command adds a different volume copy, such as a copy created from an uncompressed to compressed conversion or a compressed to uncompressed conversion.

Note: A volume cannot have volume copies in different storage pools if cloud snapshot is enabled on the volume.

An encryption key cannot be used when you add an image mode MDisk. To use encryption (when the MDisk has an encryption key), the MDisk must be self-encrypting.

Remember: You cannot add a volume copy if the volume to be copied is being formatted.

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 volume by 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 storage pool are used to create the volume. The striping is at an extent level; one extent from each managed disk in the group is used. For example, a storage pool 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 list can include two or more managed disks from the same storage pool. 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 volumes to be created when a managed disk already has data on it, perhaps from a previrtualized subsystem. When an image mode volume is created, it directly corresponds to the (previously unmanaged) managed disk that it was created from; therefore, volume logical block address (LBA) x equals managed disk LBA i . You can use this command to bring a nonvirtualized disk under the control of the system. After it is under the control of the system, you can migrate the volume from the single managed disk. When it is migrated, the volume is no longer an image mode volume.

You can add image mode volumes to an already populated storage pool with other types of volumes, such as a striped or sequential.

Note: An image mode copy must be at least as large as the volume that it is being added to, but any capacity beyond the size of the volume is not accessible.

The command returns the ID of the newly created volume copy.

Create the first compressed volume copy for an I/O group to activate compression. You cannot create or move a compressed volume copy to an I/O group that contains (at least) one node that does not support compressed volumes. You must use another I/O group, but note that this does not affect moving to the recovery I/O group.

Important:

- If the volume (or volume copy) is a target of a FlashCopy mapping with a source volume in an active-active relationship the new storage pool must be in the same site as the source volume.
- If this command is used for a volume that is a master volume, an auxiliary volume, or a change volume of an active-active relationship, the new copy must be created in a storage pool of the same site as the existing volume copy.
- When adding an image mode copy, the site information for the MDisk being added must be well-defined and match the site information for any other MDisks in the storage pool.

The rate at which the volume copies resynchronize after loss of synchronization can be specified by using the **-syncrate** parameter. Table 105 provides the relationship of the *syncrate* value to the data copied per second.

Note: These settings also affect the initial rate of formatting.

Table 105. Relationship between the *syncrate* value and the data copied per second

User-specified <i>syncrate</i> attribute value	Data copied/sec
1 - 10	128 KB
11 - 20	256 KB
21 - 30	512 KB

Table 105. Relationship between the syncrate value and the data copied per second (continued)

User-specified <i>syncrate</i> attribute value	Data copied/sec
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

```
addvdiskcopy -mdiskgrp 0 -easytier off vdisk8
```

The resulting output:

```
Vdisk [8] copy [1] successfully created
```

An invocation example for specifying storage pools

```
addvdiskcopy -mdiskgrp 0 -vtype image -mdisk 13 -tier tier0_flash -easytier off vdisk9
```

The resulting output:

```
Vdisk [9] copy [1] successfully created
```

An invocation example for configuring a mirror write algorithm priority

```
addvdiskcopy -mdiskgrp 0 -mirrorwritepriority latency vdisk9
```

The resulting output:

```
Vdisk [9] copy [1] successfully created
```

An invocation example for adding a compressed volume copy

```
addvdiskcopy -mdiskgrp 1 -rsize 10% -compressed vdisk2
```

The resulting output:

```
Vdisk [2] copy [1] successfully created
```

An invocation example for adding a compressed volume copy

```
addvdiskcopy -mdiskgrp 0 -vtype image -mdisk 13 -tier tier_nearline vdisk9
```

The resulting output:

```
Vdisk [9] copy [1] successfully created
```

addvdiskaccess

Use the **addvdiskaccess** to add an I/O group (or groups) to the set of I/O groups in which a volume can be made accessible to hosts.

Syntax

```

▶▶ addvdiskaccess — — -iogrp ——— iogrp_id_list ——— vdisk_id —————▶▶
                               | iogrp_name_list |         | vdisk_name |

```

Parameters

-iogrp *iogrp_id_list* | *iogrp_name_list*

(Required) Specifies a list of I/O groups to add to the I/O group volume access set.

vdisk_id | *vdisk_name*

(Required) Specifies the volume to which to add access through the specified I/O groups.

Description

If an I/O group is already a member of the access set, no error is generated and no action is taken for that I/O group. All host mappings for the volume are added to the I/O groups in the list. The **-force** option is not required to extend additional mappings to other I/O groups.

When an I/O group is added to the access set, it creates access to the volume from the hosts that are mapped to the volume from the nodes in the I/O group. If the volume is mapped twice, it is also mapped twice through all additional I/O groups.

You can add I/O groups to the volume access list if they are mapped to iSCSI hosts. This means that iSCSI hosts can access volumes that are accessible through multiple I/O groups (as well as single I/O groups).

Remember: The **-addvdiskaccess** command fails if:

- Any host (for which the volume has a host mapping) is not associated with an I/O group in the list
- The host volume mapping limit is exceeded
- The amount of extra mappings added exceeds the clustered system limit for host volume mappings

Two mappings are created if a host is mapped to a volume with two I/O groups. Hosts are limited to 512 host-to-volume mappings, which means a host can be mapped to:

- 512 volumes in a single I/O group
- 256 volumes across two I/O groups
- 64 volumes across four I/O groups

The command fails if any host mapped to the volume is detected as a host system that does not support volumes mapped from multiple I/O groups.

An invocation example

This example adds I/O group 2 to the volume access set for DB_Volume:

```
addvdiskaccess -iogrp 2 DB_Volume
```

The resulting output:

No feedback

An invocation example

This example adds I/O groups 2 and 3 to the volume access set for volume ID 3:

```
addvdiskaccess -iogrp 2:3 3
```

The resulting output:

No feedback

analyzevdisk

Use the **analyzevdisk** command to queue or cancel volume analysis.

Syntax

```
►► analyzevdisk — [ -cancel ] [ vdisk_id vdisk_name ] ►►
```

Parameters

-cancel

(Optional) Cancels an ongoing compression estimation.

vdisk_id | *vdisk_name*

(Required) Specifies the volume ID or name to queue for analysis.

Description

This command queues or cancels volume analysis. The order is based on the *vdisk_id* value.

Important: You cannot specify `analyzevdisk -cancel` for a volume that is not currently being analyzed (or is queued for analysis).

You can schedule an offline volume for analysis (no error message is displayed). The volume remains scheduled until it is back online and is analyzed according to its *vdisk id* value.

A concise invocation example to queue vdisk 0 for analysis

```
analyzevdisk 0
```

The detailed resulting output:

No feedback

A concise invocation example to dequeue or cancel an ongoing analysis for vdisk 0

```
analyzevdisk -cancel 0
```

The detailed resulting output:

No feedback

analyzevdiskbysystem

Use the **analyzevdiskbysystem** command to schedule all existing volumes in system for free space analysis.

Syntax

```
►► analyzevdiskbysystem — [ -cancel ] ►►
```

Parameters

-cancel

(Optional) Cancels a scheduled or pending compression estimate.

Description

This command schedules all existing volumes in system for free space analysis.

Volumes that are created after you specify the command are not evaluated. Use **analyzevdisk** to evaluate specific volumes.

A concise invocation example to queue vdisk 0 for analysis

```
analyzevdiskbysystem
```

The detailed resulting output:

No feedback

A concise invocation example to dequeue or cancel an ongoing analysis for vdisk 0

```
analyzevdiskbysystem -cancel
```

The detailed resulting output:

No feedback

backupvolume

Use the **backupvolume** command to create a volume snapshot.

Syntax

```
➤ backupvolume — [ -full ] [ volume_name | volume_id ] ➤
```

Parameters

-full

(Optional) Specifies that the snapshot generation for the volume should be a full snapshot.

volume_name | *volume_id*

(Required) Specifies the volume name or ID for the volume being backed up. The value for the volume name must be an alphanumeric string and the value for the volume ID must be a number.

Description

This command creates a volume snapshot.

The command completes when the volume snapshot is taken, and the snapshot is asynchronously transferred to the cloud system.

Note: If a volume belongs to a volume group, you must specify **backupvolumegroup** instead of **backupvolume**.

An invocation example for creating a full snapshot generation that has existing volume snapshots on the cloud system

```
backupvolume -full vdisk7
```

The resulting output:

No feedback

An invocation example for creating a backup of a volume for the first time

```
backupvolume neymar7
```

The resulting output:

No feedback

An invocation example for creating a backup of a volume that has existing snapshots in the cloud

```
backupvolume jvardy6
```

The resulting output:

No feedback

An invocation example for creating a full snapshot for a volume that has existing snapshots in the cloud

```
backupvolume -full lmessi1
```

The resulting output:

No feedback

backupvolumegroup

Use the **backupvolumegroup** command to create a new snapshot for all of the volumes in a volume group.

Syntax

```
➤ backupvolumegroup — [ -full ] [ volumegroup_name | volumegroup_id ] ➤
```

Parameters

-full

(Optional) Specifies a full backup for volume group members.

volumegroup_name | *volumegroup_id*

(Required) Specifies a volume group ID or name for the volume to back up. The value must be a number for the volume group ID and an alphanumeric string for the volume group name.

Description

This command creates a new snapshot for all of the volumes in a volume group.

The command completes as soon as a snapshot of the volume group is taken. The backup is asynchronously transferred to the cloud. If any of the volume members have a backup or restore in progress, any new volume group backup cannot be taken. The volume backup on each volume member needs to be enabled using **chvdisk** command to enable the volume backup for entire volume group.

An invocation example

To create a backup of a volume group for the first time:

```
backupvolumegroup volgroup1
```

The resulting output:

No feedback

An invocation example

To create a backup of a volume group that has existing backups in the cloud:

```
backupvolumegroup volgroup1
```

The resulting output:

No feedback

An invocation example

To create a full backup generation for a volume group that has existing backups in the cloud:

```
backupvolumegroup -full volgroup1
```

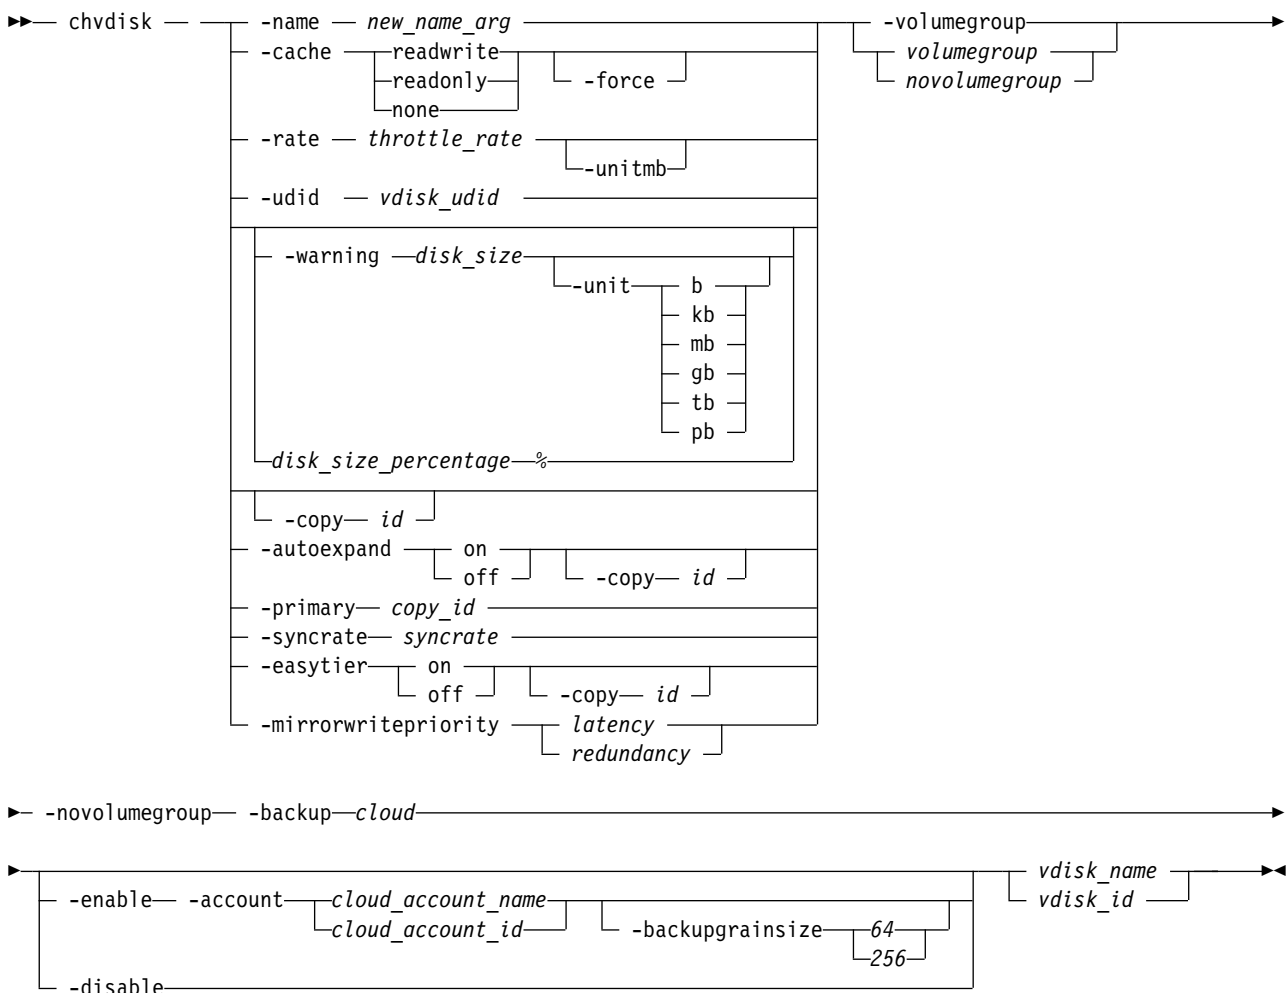
The resulting output:

No feedback

chvdisk

Use the **chvdisk** command to modify the properties of a volume, such as the disk name, I/O governing rate, or unit number. You can also change IBM Easy Tier settings.

Syntax



Parameters

-name *new_name_arg*

(Optional) Specifies a new name to assign to the volume. You cannot use this parameter with the **-rate** or **-udid** parameters. This parameter is required if you do not use the **-rate** or **-udid** parameters.

Note: Do not use this parameter with file system volumes.

-cache *readwrite* | *readonly* | *none*

(Optional) Specifies the caching options for the volume. Valid entries are:

- *readwrite* to enable the cache for the volume
- *readonly* to disable write caching while allowing read caching for a volume
- *none* to disable the cache mode for the volume

The default is *readwrite*.

Remember: If you do not specify the **-cache** parameter, the default value (*readwrite*) is used.

-force

(Optional) The **force** parameter can only be used for changing the caching mode. Use the **force** parameter with the **cache** parameter to specify that you want the system to change the cache mode of the volume even if the I/O group is offline. This option overrides the cache flush mechanism.

Attention: If the **force** parameter is used for changing the caching mode, the contents of the cache are discarded and the volume might be corrupted by the loss of the cached data. This could occur if the system is able to destage all write data from the cache or not. The **force** parameter should be used with caution.

Important: Using the force parameter might result in a loss of access. Use it only under the direction of your product support information.

-rate *throttle_rate* **-unitmb**

(Optional) Specifies the I/O governing rate for the volume, which caps the amount of I/O that is accepted. The default *throttle_rate* units are I/Os. By default the *throttle_rate* is disabled. To change the *throttle_rate* units to megabits per second (MBps), specify the **-unitmb** parameter. The governing rate for a volume can be specified by I/Os or by MBps, but not both. However, you can set the rate to I/Os for some volumes and to MBps for others. When the Input/Output Operations Per Second (IOPS) limit is configured on a volume, and it is smaller than 100 IOPS, the throttling logic rounds it to 100 IOPS. Even if throttle is set to a value smaller than 100 IOPs, the actual throttling occurs at 100 IOPs.

Note: To disable the throttling on a specific volume, set the *throttle_rate* value to zero.

You cannot use this parameter with the **-name** 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** parameters.

-warning *disk_size* | *disk_size_percentage*

(Optional) Generates a warning when the used disk capacity on the thin-provisioned 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 volume 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. The default unit value is MB.

-autoexpand on | off

(Optional) Specifies whether thin-provisioned volume copies automatically expand their real capacities by allocating new extents from their storage pool. To use this parameter, the volume must be thin-provisioned.

-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 volume is mirrored and only one volume copy is thin-provisioned. If both copies are thin-provisioned 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. You cannot change the volume's primary copy if that primary copy has its `autodelete` flag is set to `yes` (on).

Important: You cannot use this parameter with a volume that is fast formatting.

-syncrate syncrate

(Optional) Specifies the copy synchronization rate. A value of zero (0) prevents synchronization. The default value is 50. See Table 106 on page 669 for the supported **-syncrate** values and their corresponding rates. Use this parameter to alter the rate at which the fully allocated volume or mirrored volume format before synchronization.

-easytier on | off

(Optional) Enables or disables the IBM Easy Tier function.

-mirrorwritepriority latency | redundancy

(Optional) Specifies how to configure the mirror write algorithm priority. A change to the mirror write priority is reflected in the volume's view immediately and in the volume's behavior after all prior input and output (I/O) completes.

1. Choosing *latency* means a copy that is slow to respond to a write I/O becomes unsynchronized, and the write I/O completes if the other copy successfully writes the data
2. Choosing *redundancy* means a copy that is slow to respond to a write I/O synchronizes completion of the write I/O with the completion of the slower I/O in order to maintain synchronization.

-volumegroup volumegroup_name | volumegroup_id

(Optional)

Specifies a new volume group for a volume. This parameter is mutually exclusive with **-novolumegroup**.

-novolumegroup

(Optional) Specifies that a volume does not belong in any volume group. This parameter is mutually exclusive with **-volumegroup**.

-backup cloud

(Optional) Specifies the cloud snapshot type to enable or disable. The value must be `cloud`.

-enable

(Optional) Enables the backup or snapshot type specified with the **-backup** parameter.

-disable

(Optional) Disables the backup or snapshot type specified with the **-backup** parameter.

-account *cloud_account_id* | *cloud_account_name*

(Optional) Specifies the cloud account to use for the volume. You must specify **-enable** with this parameter.

-backupgrainsize *64* | *256*

(Optional) Specifies the grain size (in KB) for volume mappings. The values are 64 and 256. You must specify **-enable** to use this parameter.

You can enable a volume for a cloud snapshot with one account. You cannot enable cloud backup on a volume for a second time on the same or different cloud account.

You cannot turn off the cloud snapshot function if a snapshot in progress. Any snapshot that is in progress must complete or be canceled.

vdisk_name | *vdisk_id*

(Required) Specifies the volume to modify, either by ID or by name.

Description

The **chvdisk** command modifies a single property of a volume. To change the volume name and modify the synchronization rate, for example, you must issue the command twice. If the volume is offline, use one of the **recovervdisk** command to recover the volume and bring it back online.

Important: To change the caching I/O group for a volume or preferred node, use the **movevdisk** command.

You can specify a new name or label. You can use the new name subsequently to refer to the volume.

You can set a limit on the amount of I/O transactions that is accepted for this volume. 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 volume 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 are used. The default capacity is in MB.

When the volume is created, there is no throttling applied to it. Using the **-rate** parameter can change this. To change the volume back to an unthrottled state, specify 0 (zero) with the **-rate** parameter.

The rate at which the volume copies resynchronize after loss of synchronization can be specified by using the **-syncrate** parameter. This table provides the relationship of the *syncrate* value to the data copied per second.

Note: These settings also affect the initial rate of formatting.

Table 106. Relationship between the *syncrate* value and the data copied per second

User-specified <i>syncrate</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

Table 106. Relationship between the syncrate value and the data copied per second (continued)

User-specified <i>syncrate</i> attribute value	Data copied/sec
91 - 100	64 MB

An invocation example

```
chvdisk -rate 2040 1
```

The resulting output:

No feedback

An invocation example

```
chvdisk -cache readonly 1
```

The resulting output:

No feedback

An invocation example

```
chvdisk -volumegroup 1 vdisk2
```

The resulting output:

No feedback

An invocation example

To enable a cloud snapshot for a volume:

```
chvdisk -backup cloud -enable -account myVardyj vdisk7
```

The resulting output:

No feedback

An invocation example

To disable a cloud snapshot for a volume:

```
chvdisk -backup cloud -disable vdisk7
```

The resulting output:

No feedback

chvolumegroup

Use the **chvolumegroup** command to change volume group properties.

Syntax

```

>> chvolumegroup — [ -name — volumegroup_name ] [ volumegroup_name
                    [ volumegroup_id ] —————>>

```

Parameters

-name *volumegroup_name*

(Optional) Specifies a new volume group name. The value must be an alphanumeric string.

(Required) Specifies the volume group name or group ID for the volume you want to modify. The value must be a number for the volume group ID, and an alphanumeric string for the volume group name.

This command changes volume group properties.

```
chvolumegroup -name newname1 1
```

No feedback

Use the **expandvdisksize** command to expand the size of a volume by a given capacity.

```

>> expandvdisksize — [ -size disk_size ] [ -rsize disk_size ] [ -copy id ]
> [ -mdisk [ mdisk_id_list mdisk_name_list ] ] [ -fmtdisk ] [ -unit [ b kb mb gb tb pb ] ]
> [ vdisk_name vdisk_id ]

```

(Required) Specifies the capacity by which the volume 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 volume is thin-provisioned, MDisks cannot be specified.

(Optional) Specifies the capacity by which to increase the real size of a thin-provisioned volume. 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 volume. You cannot specify the **-rsize** parameter with the **-size** parameter. You must specify either **-size** or **-rsize**.

(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 volume copy. The **-copy** parameter is required

if the specified volume is mirrored and only one copy is thin-provisioned. If the volume is mirrored, both copies are thin-provisioned 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 volume come from the specified list of MDisks. All MDisks in the list must be part of the same storage pool. The **-mdisk** parameter cannot be used if the specified volume is mirrored.

-fmtdisk

(Optional) Specifies that the volume be formatted before use. This parameter formats the new extents that have been added to the volume as a result of the **expandvdiskspace** command. The **expandvdiskspace** 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 volume to modify, either by ID or by name.

Description

Use the **expandvdiskspace** command to expand the physical capacity that is allocated to a particular volume by the specified amount.

The command can also be used to expand the virtual capacity of a thin-provisioned volume without altering the physical capacity that is assigned to the volume. To change the capacity of a non-thin-provisioned volume, or the virtual capacity of a thin-provisioned volume, use the **-size** parameter. To change the real capacity of a thin-provisioned volume, use the **-rsize** parameter.

Note: For relationships that cannot be resized, you cannot expand the capacity of any volume in a Global Mirror or Metro Mirror relationship, regardless of whether it is a Primary, Secondary, or a Change Volume. To expand the capacity of a volume in a Global Mirror or Metro Mirror relationship:

1. Delete the relationship.
2. Increase the size of all the volumes. All volumes in a relationship must have the exact same size (virtual capacity).
3. Re-create the relationship with the larger volumes.

When the mirror is restarted, it will do a complete initial synchronization, replicating the entire primary volume to the secondary volume.

Note: You can expand the capacity of any volume in a Global Mirror or Metro Mirror relationship that is in *consistent_synchronized* state if those volumes are using thin-provisioned or compressed copies. You cannot expand the capacity for these types of volumes:

- Volumes in HyperSwap relationships or in Global Mirror relationships that are operating in *cycling* mode
- Volumes in relationships where a change volume configured
- Volumes that have a fully-allocated volume copy

You can not expand the capacity of any volume in a FlashCopy mapping, regardless of whether it is a source or target, or what state the mapping is in. You can expand the capacity of a volume in a FlashCopy mapping:

1. Delete all the mappings in that FlashCopy tree. (There is a root source volume and some targets either directly or cascaded off of other targets - the entire tree must be deleted.)

2. Increase the size of all volumes in the original FlashCopy tree. All volumes in a tree must be the same size (virtual capacity).
3. Re-create all the FlashCopy mappings with the new larger volumes.

When a FlashCopy is restarted after being deleted (including if it is an incremental FlashCopy) the entire volume becomes part of any background copy because it is the start of a new mapping.

Note: The default capacity units are MB.

When a volume is expanded, the virtualization policy can change. Its mode becomes striped even if it was previously sequential. See the **mkvdisk** command for details of the virtualization policies.

To run the **expandvdisksize** command on a mirrored volume, all copies of the volume must be synchronized. The command formats all copies of a mirrored volume automatically.

Remember:

1. You cannot resize (expand) an image mode volume.
2. You cannot resize (expand) a volume that is part of a file system.
3. You cannot resize (expand) volume if that volume is being formatted.
4. You cannot use the **-mdisk** parameter to resize (expand) a volume when a volume is being migrated.
5. You cannot use **shrinkvdisksize** if the volume is fast formatting.
6. You can not resize (expand) a volume if cloud snapshot is enabled on that volume.

You must expand both volumes in a relationship to maintain full operation of the system. To perform this:

1. Expand the secondary volume by the required additional capacity
2. Expand the primary volume by the required additional capacity

An invocation example to increase the capacity of vdisk1 by 2048 bytes using extents from two MDisks (and to format the new part of the volume)

```
expandvdisksize -size 2048 -unit b -mdisk mdisk0:mdisk1 -fmt disk 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 volume)

```
expandvdisksize -size 100 -unit mb -mdisk mdisk0:mdisk1 -fmt disk vdisk1
```

The resulting output:

No feedback

An invocation example to increase the real capacity of thin-provisioned vdisk2 by 100 MB without changing the virtual capacity (and to spread the extents across all MDisks in the storage pool)

```
expandvdisksize -rsize 100 -unit mb vdisk2
```

The resulting output:

No feedback

An invocation example to increase the real capacity of thin-provisioned volume copy id 1 of mirrored volume vdisk3 by 100 MB

```
expandvdisksize -rsize 100 -unit mb -copy 1 vdisk3
```

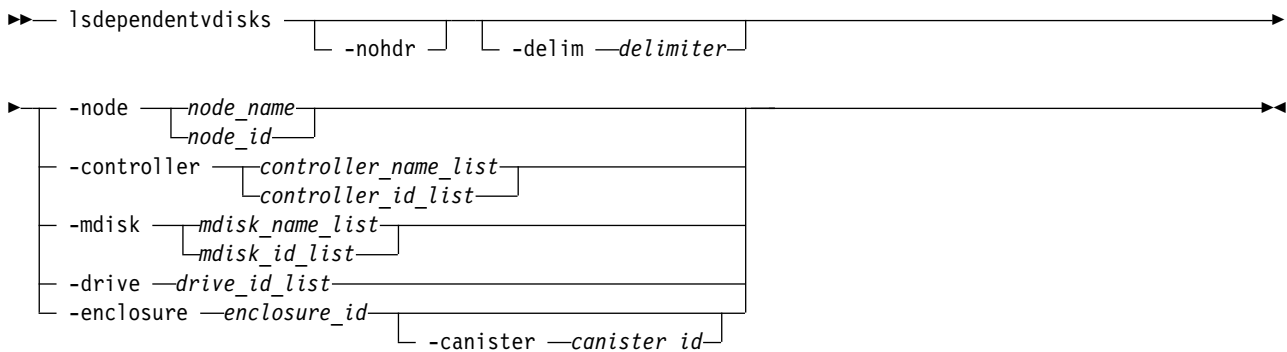
The resulting output:

No feedback

Isdependentvdisks

Use the **Isdependentvdisks** command to view which volumes go offline if you remove a specific piece of hardware from the system.

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.

-node node_name | node_id

(Optional) Specifies the node for which volume dependency is required.

-controller controller_name_list | controller_id_list

(Optional) Specifies the controllers for which volume dependency is required.

-mdisk mkdisk_name_list | mkdisk_id_list

(Optional) Specifies the MDisks for which volume dependency is required.

-drive

(Optional) Specifies the drives for which volume dependency is required. There is a maximum of 128 entries.

-enclosure *enclosure_id*

(Optional) Specifies the enclosure for which volume dependency is required. You can remove a control enclosure without affecting your other data.

-canister *canister_id*

(Optional) Specifies an enclosure canister. The possible values are 1 and 2.

Description

Use this command to view which volumes go offline if you remove a specific piece of hardware from the system. Use this command to determine which volumes are affected before undergoing maintenance.

An invocation example

```
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 go offline.

lshostvdiskmap

Use the **lshostvdiskmap** command to display a list of volumes mapped to a given host. These are the volumes that are recognized by the specified host.

Syntax

```
➤ lshostvdiskmap [ -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*

(Optional) Specifies the host in terms of its ID or name. The command displays a list of all the volumes that are mapped to the specified host and the Small Computer System Interface (SCSI) ID by which they are mapped. If neither a host ID or name are entered, the command displays a list of all recognized volume mappings.

Description

This command displays a list of volume IDs and names. These are the volumes 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 volume is recognized by the host.

Each volume that is exported by the clustered system (system) is assigned a unique virtual path (VPATH) number. This number identifies the volume and determines which volume corresponds to the volume that the hosts recognize. This procedure can only be completed using the command-line interface.

For a specific volume based on which operating system and multipath software are used, you can use different commands to determine the VPATH serial number. For example, issuing **datapath query device** finds the VPATH serial number for volumes mapped to AIX `sddpcm`.

Find the host that is defined to the system that corresponds with the host that you are working with.

1. The worldwide port names (WWPNs) are an attribute of the host bus adapter (HBA). You can find these by looking at the device definitions stored by your operating system. For example, on AIX® they are in the Object Data Manager (ODM), in Windows® they are in the Device Manager details for the given HBA.
2. Verify which host is defined to the system 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:

```
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 system orange.

When you have the *hostname* defined to the system and the *vpath serial number*, issue the following command:

```
lshostvdiskmap hostname
```

where *hostname* is the name of the host. A list is displayed. Look for the volume UID that matches the *vpath serial number* and record the volume name or ID.

The command returns the following values:

id Indicates the host ID in the output for **lshostvdiskmap**.

name Indicates the host name in the output for **lshostvdiskmap**.

SCSI_id
Specifies the SCSI ID.

host_cluster_id
Indicates the unique ID for a host system.

host_cluster_name
Indicates the unique name for a host system.

vdisk_id
Indicates the ID of the volume.

vdisk_name
Indicates the name of the volume.

vdisk_UID
Indicates the UID of the volume.

IO_group_id

Indicates the ID of the input/output (I/O) group in which the host volume mapping exists.

IO_group_name

Specifies the name of I/O group in which the host volume mapping exists.

An invocation example

```
lshostvdiskmap -delim : 2
```

The resulting output:

```
id:name:SCSI_id:host_id:host_name:vdisk_id:vdisk_name:vdisk_UID:IO_group_id:IO_group_name
2:host2:0:5:var10:vdisk10:600507680195800150000000000000A:0:iogrp0
2:host2:1:4:var2:vdisk11:600507680195800150000000000000B:1:iogrp1
2:host2:2:3:var3:vdisk12:600507680195800150000000000000C:0:iogrp0
2:host2:3:2:var4:vdisk13:600507680195800150000000000000D:1:iogrp1
2:host2:4:1:var5:vdisk14:600507680195800150000000000000E:1:iogrp0
```

An invocation example

```
lshostvdiskmap 0
```

The resulting output:

id	name	SCSI_id	host_id	host_name	vdisk_UID	IO_group_id	IO_group_name	mapping_type	host_cluster
0	vdisk0	0	1	hvlab02c2	6005076801D901A3F800000000000000	0	io_grp0	shared	0
0	vdisk0	0	4	vmlab02c1	6005076801D901A3F800000000000000	0	io_grp0	shared	0
0	vdisk0	0	5	vmlab02c2	6005076801D901A3F800000000000000	0	io_grp0	shared	0
0	vdisk0	0	24	vmlab14c1	6005076801D901A3F800000000000000	0	io_grp0	shared	0
0	vdisk0	0	25	vmlab14c2	6005076801D901A3F800000000000000	0	io_grp0	shared	0
0	vdisk0	0	26	vmlab15	6005076801D901A3F800000000000000	0	io_grp0	private	

lsmetadatavdisk

Use the **lsmetadatavdisk** command to display the information for metadata volume.

Syntax

```
➤ lsmetadatavdisk — [ -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 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 1-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 information for metadata volume.

Table 107 provides the attribute values that can be displayed as output view data.

Table 107. Ismetadatavdisk output

Attribute	Description
vdisk_od	Indicates the ID of the metadata volume.
vdisk_name	Indicates the name of the metadata volume.
status	Indicates the running status of the metadata volume.

An invocation example

To display information for a metadata volume:

```
lsmetadatavdisk
```


The resulting output:

```
vdisk_id      2
vdisk_name    vdisk2
status        online
```

Isrepairsevdiskcopyprogress

The **lsrepairsevdiskcopyprogress** command lists the repair progress for thin-provisioned volume copies or compressed volume 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 volume 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 thin-provisioned copies in the clustered system.

Description

The **lsrepairsevdiskcopyprogress** command lists the repair progress for thin-provisioned or compressed copies for the specified volume. If you do not specify a volume, the command lists the repair progress for all thin-provisioned or compressed copies in the clustered system.

Remember: Only run this command after running the **repairsevdiskcopy** command, which you must only run as required by the fix procedures or by your product support information.

The command returns values for the following volume copy attributes:

task Specifies the active task.

- repairing indicates repair of a thin-provisioned volume copy
- compressed_repairing indicates repair of a compressed volume copy.

progress

Specifies the task completion percentage.

estimated_completion_time

Specifies the expected duration of the task in the format YYMMDDHHMMSS (or blank if the estimated completion time is unknown).

An invocation example

```
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

```
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

```
lsrepairsevdiskcopyprogress -delim : -copy 1 vdisk0
```


The resulting output:

```
id:name:copy id:task:progress:estimated_completion_time
0:vdisk0:1:repairing:51:070301120000
```

lsrepairvdiskcopyprogress

Use the **lsrepairvdiskcopyprogress** command to display the progress of 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 volume 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 volume 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 volumes. Use this command to track progress after running the **repairvdiskcopy** command. You can specify a volume copy using the **-copy** parameter. To display the volumes that have two or more copies with an active task, specify the command with no parameters; it is not possible to have only one volume copy with an active task.

The command displays progress for the following types of volume copies:

- All volume copies display the same task; validate, medium or resync, depending on the specified parameter.
- All volume copies display the same percentage and estimated completion time.
- If specified, non-mirrored volumes 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.

The command returns values for the following volume repair attributes:

vdisk_id

Indicates the volume ID.

vdisk_name

Indicates the volume name.

copy_id

Indicates the system-assigned identifier for the volume copy.

task Indicates the active task. The values can be `repairing` or `compressed_repairing`.

progress

Indicates the task completion percentage. This value is 0 when task is in `compressed_repairing` state.

estimated_completion_time

Indicates the expected time (duration) the task completion time. The value is in the `YYMMDDHHMMSS` format, and is blank if the duration is not known.

An invocation example

```
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

```
lsrepairvdiskcopyprogress -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

```
lsrepairvdiskcopyprogress -delim : -copy 0 vdisk0
```

The resulting output:

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

An invocation example showing a compressed volume copy and a TP volume copy being repaired

```
lsrepairvdiskcopyprogress
```

The resulting output:

vdisk_id	vdisk_name	copy_id	task	progress	estimated_completion_time
0	vdisk0	0	repairing	50	070301120000
2	vdisk2	1	compressed_repairing	0	070301080102

lssevdiskcopy


Use the **lssevdiskcopy** command to list the thin-provisioned copies of the specified volumes.

Syntax

```

▶▶—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 volume copy to list thin-provisioned 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 **lssevdiskcopy** command are valid:

- *mdisk_grp_id*
- *mdisk_grp_name*
- *overallocation*
- *autoexpand*
- *grainsize*

vdisk_name | *vdisk_id*

(Optional) Specifies the volume name or ID to list thin-provisioned copies for. You must specify this parameter last on the command line. If you do not enter this parameter, the command lists all thin-provisioned copies in the clustered system.

Description

The **lssevdiskcopy** command lists all thin-provisioned copies of the specified volume. If you do not specify a volume, the command lists all thin-provisioned volume copies in the clustered system.

The command provides a concise view of the thin-provisioned properties of the selected volume copies. Run the **lsvdiskcopy** command to see a concise view of the properties that are common to thin-provisioned and non-thin-provisioned volume copies. See the description of the **lsvdisk** command for a description of the fields that are shown in the view.

The command returns values for the following volume copy attributes:

copy_id Indicates a system-assigned identifier for the volume copy. The value can be 0 or 1.

status Indicates the system status. The value can be online or offline. A copy is offline if all nodes cannot access the storage pool that contains the copy.

sync Indicates whether the volume copy is synchronized.

auto_delete Indicates the primary copy is deleted once the secondary copy is synchronized. The values are yes or no.

primary Indicates whether the volume copy is the primary copy. A volume has exactly one primary copy. The value can be yes or no.

mdiskgrp_id Indicates the ID of the storage pool that the volume copy belongs to.

mdiskgrp_name Indicates the name of the storage pool that the volume copy belongs to.

type Indicates the virtualization type of the volume. The value can be striped, sequential or image.

mdisk_id Indicates the MDisk ID that is used for sequential and image mode volumes.

mdisk_name Indicates the MDisk name that is used for sequential and image mode volumes.

fast_write_state Indicates the cache state of the volume copy. The value can be empty, not_empty, corrupt, or repairing. The value is always empty for non-thin-provisioned copies. A cache state of corrupt indicates that the volume is thin-provisioned and requires repair that is initiated by a **recovervdisk** command or the **repairsevdiskcopy** command.

used_capacity Indicates the portion of **real_capacity** that is being used to store data. For non-thin-provisioned copies, this value is the same as the volume capacity. If the volume copy is thin-provisioned, the value increases from zero to the **real_capacity** value as more of the volume is written to.

real_capacity Indicates the amount of physical storage that is allocated from an storage pool to this volume copy. If the volume copy is not thin-provisioned, the value is the same as the volume capacity. If the volume copy is thin-provisioned, the value can be different.

free_capacity Indicates the difference between the **real_capacity** and **used_capacity** values.

overalllocation Expressed as a percentage, indicates the ratio of volume capacity to **real_capacity** values. This value is always 100 for non-thin-provisioned volumes.

autoexpand Indicates whether autoexpand is enabled on a thin-provisioned volume. The value can be on or off.

warning Expressed as a percentage, for thin-provisioned volume copies only, indicates that a warning is generated when the ratio of **used_capacity** to volume capacity reaches the specified level.

grainsize For thin-provisioned volume copies, indicates the grain size chosen for the volume copy when it was created.

Remember: This field is always blank for compressed volume copies.

se_copy

Indicates if the copy is thin-provisioned. The value can be yes or no.

easy_tier

Indicates 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

Indicates which Easy Tier functions are active for the volume copy:

- active indicates a pool is being managed by Easy Tier to provide tier management performance-based pool balancing. For example, extents of this volume copy can be moved for performance (automatic data placement).
- inactive indicates that no Easy Tier function is active.
- measured indicates that statistics are being gathered for this volume copy, but no extents are moved.
- balanced indicates a pool is being managed by Easy Tier to provide performance-based pool balancing (for example, extents can be moved).

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 685)
Off	One	On	inactive (see note 2 on page 685)
Off	Two	Off	inactive (see note 2 on page 685)
Off	Two	On	inactive (see note 2 on page 685)
Measure	One	Off	measured (see note 3 on page 685)
Measure	One	On	measured (see note 3 on page 685)
Measure	Two	Off	measured (see note 3 on page 685)
Measure	Two	On	measured (see note 3 on page 685)
Auto	One	Off	measured (see note 3 on page 685)
Auto	One	On	balanced (see note 4 on page 685)
Auto	Two	Off	active (see note 5 on page 685)
Auto	Two	On	measured (see note 3 on page 685)
On	One	Off	measured (see note 3 on page 685)

Storage pool Easy Tier setting	Number of tiers in the storage pool	Volume copy Easy Tier setting	Volume copy Easy Tier status
On	One	On	balanced (see note 4)
On	Two	Off	measured (see note 3)
On	Two	On	active (see note 5)

Note:

1. If the volume copy is in image or sequential mode or is being migrated then the volume copy Easy Tier(tm) status is measured instead of active.
2. When the volume copy status is inactive, no Easy Tier(tm) functions are enabled for that volume copy.
3. When the volume copy status is measured, the Easy Tier(tm) function collects usage statistics for the volume but automatic data placement is not active.
4. When the volume copy status is balanced, the Easy Tier(tm) function enables performance-based pool balancing for that volume copy.
5. When the volume copy status is active, the Easy Tier(tm) function operates in automatic data placement mode for that volume.
6. The default Easy Tier(tm) setting for a storage pool is auto, and the default Easy Tier(tm) setting for a volume copy is on. This means that Easy Tier(tm) functions except pool performance balancing are disabled for storage pools with a single tier, and that automatic data placement mode are enabled for all striped volume copies in a storage pool with two or more tiers.

tier Indicates which tier information is being reported:

- tier0_flash
- tier1_flash
- tier_enterprise
- tier_nearline

tier_capacity

Indicates the total MDisk capacity assigned to the volume in the tier.

Note: For thin-provisioned copies, the capacity by tier will be the real capacity.

compressed_copy

Indicates whether or not the volume copy is compressed.

uncompressed_used_capacity

For compressed volumes, indicates the amount of data written to the volume before compression.

parent_mdisk_grp_id

Indicates the physical storage pool id that the volume extents are actually allocated from. This is a numeric string (between 0 and 127 characters) or blank.

parent_mdisk_grp_name

Indicates the physical storage pool name that the volume extents are actually allocated from. This is an alphanumeric string (between 1 and 63 characters) or blank.

encrypt

Indicates whether the volume and its copies are encrypted. The values are yes or no.

An invocation example

`lssevdiskcopy -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:se_copy:compressed_copy
:uncompressed_used_capacity
```

```
0:vv1:0:0:ppp:16.00GB:2.00GB:2.01GB:6.00GB:796:off:20:32:no:yes:3.27GB
1:se1:0:0:ppp:16.00GB:1.00GB:4.00GB:15.00GB:400:off:20:32:yes:no:1.0GB:yes:no:1.0GB
1:se1:1:0:ppp:16.00GB:2.00GB:2.01GB:14.00GB:796:off:45:256:no:yes:4.46GB
```

An invocation example

```
lssevdiskcopy -delim : -copy 0 0
```

The resulting output:

```
vdisk_id:0
vdisk_name:vv1
capacity:16.00GB
copy_id:0
status:online
sync:yes
auto_delete: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:2.01GB
free_capacity:6.00GB
overallocation:796
autoexpand:on
warning:25
grainsize:256
se_copy:yes
easy_tier:on
easy_tier_status:active

tier:tier0_flash
tier_capacity:1.63TB
tier:tier1_flash
tier_capacity:1.63TB
tier:tier_enterprise
tier_capacity:0.00MB
tier:tier_nearline
tier_capacity:0.00MB
tier_capacity:64.00MB
tier:ssd
tier_capacity:2.00GB
compressed_copy:yes
uncompressed_used_capacity:3.27GB

parent_mdisk_grp_id:10
parent_mdisk_grp_name:pool10
encrypt:yes
```

An invocation example

```
lssevdiskcopy -copy 0 -delim : vv1
```

The resulting output:

```
vdisk_id:0
vdisk_name:vv1
capacity:16.00GB
copy_id:0
status:online
sync:yes
auto_delete:yes
primary:yes
mdisk_grp_id:1
```



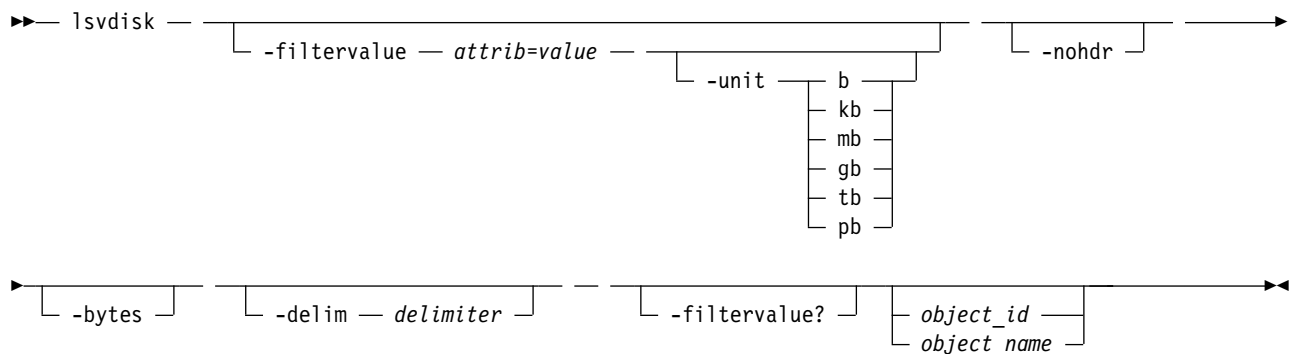
```
mdisk_grp_name:mdisk_group_1
type:striped
mdisk_id:
mdisk_name:
fast_write_state:empty
used_capacity:2.00GB
real_capacity:8.00GB
free_capacity:6.00GB
overallocation:200
autoexpand:on
warning:25
grainsize:256
se_copy:yes
easy_tier:off
easy_tier_status:inactive

tier:tier0_flash
tier_capacity:1.63TB
tier:tier1_flash
tier_capacity:1.63TB
tier:tier_enterprise
tier_capacity:0.00MB
tier:tier_nearline
tier_capacity:0.00MB
compressed_copy:no
uncompressed_used_capaity:8.00GB
parent_mdisk_grp_id:10
parent_mdisk_grp_name:pool10
encrypt:yes
```

lsvdisk

Use the **lsvdisk** command to display a concise list or a detailed view of volumes that are recognized by the clustered system (system).

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 displayed. If a capacity is specified, the units must also be included.

-unit **b** | **kb** | **mb** | **gb** | **tb** | **pb**

(Optional) Specifies the data units for the **-filtervalue** parameter.

Note: **-unit** must be used with **-filtervalue**.

-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 thin-provisioned 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.

-filtervalue?

(Optional) Displays a list of valid filter attributes. The following filters for the **lsvdisk** command are valid:

- access_IO_group_count
- backup_status
- capacity
- cloud_backup_enabled
- cloud_account_id
- cloud_account_name
- compressed_copy_count
- copy_count
- fast_write_state
- FC_id
- fc_map_count
- FC_name
- filesystem
- function
- id
- IO_group_id
- IO_group_name
- mdisk_grp_name
- mdisk_grp_id
- mirror_write_priority
- name
- owner_type
- owner_id
- owner_name
- preferred_node_id
- RC_change
- RC_id

- RC_name
- restore_status
- se_copy_count
- status
- type
- vdisk_UID
- volume_group_id
- volume_group_name
- volume_id
- volume_name

Note: It is not possible to filter the **lsvdisk** command with `mdisk_grp_name=many` to identify mirrored volumes. Instead, filter on `copy_count=2`.

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

This command displays a concise list or a detailed view of attributes for all volumes and volume copies in the system.

The volume 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 volume.
- All synchronized copies for this volume are in storage pools that are offline.
- The volume is formatting.

If you have a degraded volume and all of the associated nodes and MDisks are online, refer to your product support information for assistance. A volume 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 MDisks in the storage pool that the volume spans. In this case MDisks are shown as degraded and the fix procedures for MDisks should be followed to resolve the problem.
- The fast write cache pins data for one or more volumes 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 fix procedures for this event log to resolve the problem. The most common causes of pinned data are the following:
 - One or more volumes 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 fabric faults or misconfiguration, back-end controller faults or misconfiguration or because repeated errors has led to the system excluding access to a MDisk through one or more nodes.
 - One or more volumes in an I/O group is offline due to a problem with a FlashCopy mapping.
 - A thin-provisioned disk has run out of space.

You can encrypt volumes and volume copies. The volume is encrypted if all volume copies are also encrypted.

Note: This means that during a migration of a volume (with one copy) between an encrypted and unencrypted storage pool the value is no.
The volume is not encrypted if the storage pool has a value of `encrypt:no`, including if the volume's extents are part of an encrypted MDisk or sequential volume.

The command returns values for the following volume attributes:

IO_groups_id

Indicates the I/O group (ID) that the volume belongs to.

IO_groups_name

Indicates the I/O group (name) that the volume belongs to.

status Indicates the status. The value can be `online`, `offline` or `degraded`.

For online HyperSwap volumes the scope of offline copy is included with the status information. For offline volumes the auxiliary copy is included with the status information even if the copy associated with that is online.

mdisk_grp_id

Indicates the ID of the storage pool that the volume belongs to. If the volume has more than one copy, these fields display many.

mdisk_grp_name

Indicates the name of the storage pool that the volume belongs to. If the volume has more than one copy, these fields display many.

type Indicates the virtualization type of the volume. The value can be `striped`, `seq`, `image` or `many`. The value `many` indicates that the volume has more than one copy, which can have different virtualization types.

capacity

Indicates the total capacity of the volume.

formatted

Indicates whether the volume was formatted when it was created. The value can be `yes` or `no`.

formatting

Indicates whether the volume is formatting. The value can be `yes` or `no`.

mdisk_id

Indicates the MDisk ID that is used for sequential and image mode volume s. If the volume has more than one copy, these fields display many.

mdisk_name

Indicates the MDisk name that is used for sequential and image mode volumes. If the volume has more than one copy, these fields display many.

FC_id Indicates the ID of the FlashCopy mapping that the volume belongs to. The value `many` indicates that the volume belongs to more than one FlashCopy mapping.

FC_name

Indicates the name of the FlashCopy mapping that the volume belongs to. The value `many` indicates that the volume belongs to more than one FlashCopy mapping.

RC_id Indicates the ID of the remote copy relationship that the volume belongs to. The value must be numerical.

RC_name

Indicates the name of the remote copy relationship that the volume belongs to.

vdisk_UID

Indicates the UID of the volume.

throttle_ID

Indicates the ID for the throttle object. The value is a numeric string from 0 to 10241 (or is blank if no throttle is configured).

throttle_name

Indicates the name of the throttle object. The value is an alphanumeric string 1 to 63 characters or blank if no throttle is specified.

IOPs_limit

Indicates the IOPs limit that is configured for the volume. The value is a numeric string or blank if no limit is specified.

bandwidth_limit_MB

Indicates the bandwidth limit configured (in MB) for the volume. The value is a numeric string or is blank if no limit is configured.

preferred_node_id

Indicates the node processing the I/O data.

Remember: This value must be numeric. (The value is zero if no node is configured in the I/O group that contains the preferred node.)

fast_write_state

Indicates the cache state for the volume or volume copy. The value can be empty, not_empty, corrupt, or repairing. A cache state of corrupt indicates that the volume or volume copy requires repairing or recovery by using either the **recovervdisk** or **repairvdiskcopy** command.

cache Indicates the cache mode of the volume. The value can be readonly, readwrite, or none.

udid Indicates the unit number for the volume. Only OpenVMS hosts require a unit number.

fc_map_count

Indicates the number of FlashCopy mappings that the volume belongs to.

sync_rate

Indicates the rate for synchronization for mirrored copies.

se_copy_count

Indicates the number of thin-provisioned copies.

Remember: This value represents only thin-provisioned copies and is not used for compressed volume copies.

filesystem

Expressed as a value string (long object name with a maximum of 63 characters), indicates the full name for file system which owns this volume; otherwise, it is blank.

mirror_write_priority

Indicates the mirror write algorithm priority being used if the volume is mirrored.

RC_change

Indicates if a volume is a change volume of a remote copy relationship.

compressed_copy_count

Indicates the number of compressed volume copies.

access_IO_group_count

Indicates the number of I/O groups in the volume access set.

The command returns values for the following volume copy attributes:

copy_id

Indicates a system-assigned identifier for the volume copy. The value can be 0 or 1.

status Indicates the status. The value can be online or offline. A copy is offline if all nodes cannot access the storage pool that contains the copy.

sync Indicates whether the volume copy is synchronized.

auto_delete

Indicates the primary copy is deleted once the secondary copy is synchronized. The values are yes or no.

primary

Indicates whether the volume copy is the primary copy. A volume has exactly one primary copy. The value can be Yes or No.

mdiskgrp_id

Indicates the ID of the storage pool that the volume copy belongs to.

mdiskgrp_name

Indicates the name of the storage pool that the volume copy belongs to.

type Indicates the virtualization type of the volume. The value can be striped, seq, or image.

mdisk_id

Indicates the MDisk ID that is used for sequential and image mode volumes.

mdisk_name

Indicates the MDisk name that is used for sequential and image mode volumes.

used_capacity

Indicates the portion of `real_capacity` that is being used to store data. For non-thin-provisioned copies, this value is the same as the volume capacity. If the volume copy is thin-provisioned, the value increases from zero to the `real_capacity` value as more of the volume is written to.

real_capacity

Indicates the amount of physical storage that is allocated from an storage pool to this volume copy. If the volume copy is not thin-provisioned, the value is the same as the volume capacity. If the volume copy is thin-provisioned, the value can be different.

free_capacity

Indicates the difference between the `real_capacity` and `used_capacity` values.

overalllocation

Expressed as a percentage of the volume capacity, indicates the ratio of volume capacity to `real_capacity` values. This value is always 100 for non-thin-provisioned or compressed volumes.

Remember: This value can be any percentage (but not blank) for compressed volume copies.

autoexpand

Indicates whether **autoexpand** is enabled on a thin-provisioned volume. The value can be on or off.

Remember: This value cannot be blank for compressed copies.

warning

Expressed as a percentage of the volume capacity, this indicates a warning for thin-provisioned or compressed volume copies. A warning is generated when the ratio of `used_capacity` to volume capacity reaches the specified level.

Remember: This value can be any percentage for compressed volume copies.

grainsize

For thin-provisioned volume copies, indicates the grain size chosen for the volume copy when it was created.

Remember: This value is always blank for compressed volume copies.

se_copy

Indicates if the copy is thin-provisioned.

Remember: This value is **yes** for thin-provisioned copies and **no** for compressed volume copies.

easy_tier

This value is set by the user and indicates 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

Indicates which Easy Tier functions are active for the volume copy:

- **active** indicates a pool is being managed by Easy Tier to provide tier management performance-based pool balancing. For example, extents of this volume copy can be moved for performance (automatic data placement).
- **inactive** indicates that no Easy Tier function is active.
- **measured** indicates that statistics are being gathered for this volume copy, but no extents are moved.
- **balanced** indicates a pool is being managed by Easy Tier to provide performance-based pool balancing (for example, extents can be moved).

This table displays possible values and related information for `easy_tier_status`:

Table 108. Easy Tier status values. Easy Tier status values

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 694)
Off	One	On	inactive (see note 2 on page 694)
Off	Two	Off	inactive (see note 2 on page 694)
Off	Two	On	inactive (see note 2 on page 694)
Measure	One	Off	measured (see note 3 on page 694)
Measure	One	On	measured (see note 3 on page 694)
Measure	Two	Off	measured (see note 3 on page 694)
Measure	Two	On	measured (see note 3 on page 694)
Auto	One	Off	measured (see note 3 on page 694)
Auto	One	On	balanced (see note 4 on page 694)
Auto	Two	Off	measured (see note 3 on page 694)
Auto	Two	On	(see note 5 on page 694)

Table 108. Easy Tier status values (continued). Easy Tier status values

Storage pool Easy Tier setting	Number of tiers in the storage pool	Volume copy Easy Tier setting	Volume copy Easy Tier status
On	One	Off	measured (see note 3)
On	One	On	balanced (see note 4)
On	Two	Off	measured (see note 3)
On	Two	On	active (see note 5)

Note:

1. If the volume copy is in image or sequential mode or is being migrated then the volume copy Easy Tier(tm) status is measured instead of active.
2. When the volume copy status is inactive, no Easy Tier(tm) functions are enabled for that volume copy.
3. When the volume copy status is measured, the Easy Tier(tm) function collects usage statistics for the volume but automatic data placement is not active.
4. When the volume copy status is balanced, the Easy Tier(tm) function enables performance-based pool balancing for that volume copy.
5. When the volume copy status is active, the Easy Tier(tm) function operates in automatic data placement mode for that volume.
6. The default Easy Tier(tm) setting for a storage pool is auto, and the default Easy Tier(tm) setting for a volume copy is on. This means that Easy Tier(tm) functions except pool performance balancing are disabled for storage pools with a single tier, and that automatic data placement mode are enabled for all striped volume copies in a storage pool with two or more tiers.

tier The tier information being reported:

- tier0_flash
- tier1_flash
- tier_enterprise
- tier_nearline

tier_capacity

The total MDisk capacity assigned to the volume in the tier.

Note: For thin-provisioned copies, the capacity by tier will be the real capacity.

compressed_copy

Indicates if the volume copy is compressed.

uncompressed_used_capacity

For compressed volumes, indicates the amount of data written to the volume before compression.

last_access_time

Indicates the time (YYMMDDHHMMSS) the volume last received Small Computer System Interface (SCSI) commands from any mapped host.

parent_mdisk_grp_id

Indicates the physical storage pool id that the volume extents are actually allocated from. This is a numeric string (between 0 and 127 characters) or blank.

parent_mdisk_grp_name

Indicates the physical storage pool name that the volume extents are actually allocated from. This is an alphanumeric string (between 1 and 63 characters) or blank.

owner_type

Indicates the type of owning component or object (such as a file system). The values are:

- filesystem
- host_integration_metadata

- vvol
- none

The value is none if no owner is assigned.

owner_id

Indicates an identification number for the owning object. This value must be a numeric character unless there is no owning object (in which case it is blank).

owner_name

Indicates the name for owning object that owns this volume. The value must be a set of up to 63 alphanumeric characters, but is blank if there is no owning object.

encrypt

Indicates whether all copies of a volume reside in MDisk groups (storage pools) that are reported as encrypting. This means that either one of the following things apply:

- The copies reside in a storage pool that has an encryption key.
- All MDisk in the group are self-encrypting or encrypted for RAID.

The values are yes and no.

volume_id

Indicates the volume ID (for a high availability volume). This must be a numerical value. For a basic or stretched volume, volume_ID has the same value as id. For a HyperSwap volume (involved in an active-active relationship), the volume ID is the same as the master volume.

volume_name

Indicates the volume name (for a high availability volume). This must be an alphanumeric string containing up to 63 characters. For a basic or stretched volume, volume_name has the same value as name. For a HyperSwap volume (involved in an active-active relationship), the volume name is the same as the master volume.

function

Indicates the function of the volume in the remote copy relationship. Remote copy includes Metro Mirror, Global Mirror, and HyperSwap.

The values are:

- master, which indicates a master volume in a remote copy relationship.
- aux, which indicates an auxiliary volume in a remote copy relationship.
- master_change, which indicates a change volume for a master volume in a remote copy relationship.
- aux_change, which indicates a change volume for an auxiliary volume in a remote copy relationship.
- Blank, which indicates that the volume is not in any remote copy relationship

volume_group_id

Indicates a volume group ID for a volume group that a volume belongs to. The value must be a number.

volume_group_name

Indicates the volume group name for a volume group that a volume belongs to. The value must be an alphanumeric string.

cloud_backup_enabled

Indicates whether the cloud snapshot feature is enabled for the specified volume. The values are yes or no.

cloud_account_id

Indicates the cloud account ID. The value must be a number.

cloud_account_name

Indicates the cloud account name. The value must be an alphanumeric string.

backup_status

Indicates whether a new cloud snapshot can be started. If a backup is in progress, the status of the backup operation is given. The values are:

- off
- ready
- copying
- copying_error
- not_ready

last_backup_time

Indicates the time of the most recent backup or snapshot for the specified volume. The value must be in YYMMDDHHMMSS format (or blank).

restore_status

Indicates whether a restore can be performed for the volume. If a restore is in progress, the status of the restore operation is given. The values are:

- none
- available
- restoring
- restoring_error
- committed
- committing
- committing_error

backup_grain_size

Indicates the grain size for the volume mappings used for the cloud snapshot function. The value is blank if cloud snapshot is not enabled.

Note: This does not reflect the size of the grains stored in the cloud (which are fixed at 256 KB).

A detailed invocation example for a volume

```
lsvdisk -delim : vv45
```

The resulting output:

```
name:vv45
IO_group_id:0
IO_group_name:io_grp0
status:online
mdisk_grp_id:0
mdisk_grp_name:Group0
capacity:1000.00MB
type:striped
formatted:no
formatting:yes
mdisk_id:
mdisk_name:
FC_id:
FC_name:
RC_id:
RC_name:
vdisk_UID:60050768019B82328000000000000010

preferred_node_id:2
fast_write_state:empty
cache:readwrite
```

```

udid:
fc_map_count:0
sync_rate:50
copy_count:1
se_copy_count:0
filesystem:
mirror_write_priority:redundancy
RC_change:no
compressed_copy_count:0
access_IO_group_count:1
parent_mdisk_grp_id:5
parent_mdisk_grp_name:p5
encrypt:yes
volume_id:0
volume_name:homer0
function:aux
owner_type filesystem
owner_id 2
owner_name myfilesystem2

copy_id:0
status:online
sync:yes
auto_delete:yes
primary:yes
mdisk_grp_id:0
mdisk_grp_name:Group0
type:striped
mdisk_id:
mdisk_name:
fast_write_state:empty
used_capacity:1000.00MB
real_capacity:1000.00MB
free_capacity:0.00MB
overallocation:100
autoexpand:
warning:
grainsize:
se_copy:no
easy_tier:on
easy_tier_status:inactive

tier:tier0_flash
tier_capacity:1.63TB
tier:tier1_flash
tier_capacity:1.63TB
tier:tier_enterprise
tier_capacity:0.00MB
tier:tier_nearline
tier_capacity:0.00MB
compressed_copy:no
uncompressed_used_capacity:1000.00MB
last_access_time:140604171325
throttle_id:1
throttle_name:1cyfoxes_1
IOPs_limit:25000
bandwidth_limit_MB:500
volume_group_id:1
volume_group_name:Z1aIbra2
cloud_backup_enabled:no
cloud_account_id:
cloud_account_name:
backup_status:off
last_backup_time:
restore_status:available
backup_grain_size:

```

A concise invocation example

```
lsvdisk -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:RC_change:compressed_copy_count:volume_id:volume_name:functio
0:vdisk0:0:io_grp0:degraded:0:mdiskgrp0:10.00GB:striped:::::600507680I8300003000000000000000:0:1:empty:0:no:0:1:VDisk1:aux_
```

A detailed invocation example

```
lsvdisk -delim : vv1
```

The resulting output:

```
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
formatting:yes
mdisk_id:many
mdisk_name:many
FC_id:
FC_name:
RC_id:
RC_name:
vdisk_UID:0000000000000000AB:6005076801CF003F2800000000000000
```

```
preferred_node_id:1
fast_write_state:empty
cache:readwrite
udid:1234
fcmap_count:0
sync_rate:25
copy_count:2
se_copy_count:1filesystem:
mirror_write_priority:redundancy
RC_change:no
compressed_copy_count:0
access_IO_group_count:1
parent_mdisk_grp_id:5
parent_mdisk_grp_name:p5
encrypt:yes
volume_id:1
volume_name:slayer1
function:aux
owner_type filesystem
owner_id 2
owner_name myfilesystem2
```

```
copy_id:0
status:online
sync:yes
auto_delete:yes
primary:yes
mdisk_grp:1
mdisk_grp_name:mdisk_group_1
type:striped
mdisk_id:
mdisk_name:
```

```

fast_write_state:corrupt
used_capacity:8.00GB
real_capacity:8.00GB
free_capacity:6.00GB
overallocation:100
autoexpand:off
warning:
grainsize:
se_copy:no
easy_tier:off
easy_tier_status:inactive

tier:tier0_flash
tier_capacity:1.63TB
tier:tier1_flash
tier_capacity:1.63TB
tier:tier_enterprise
tier_capacity:0.00MB
tier:tier_nearline
tier_capacity:0.00MB

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:tier0_flash
tier_capacity:1.63TB
tier:tier1_flash
tier_capacity:1.63TB
tier:tier_enterprise
tier_capacity:0.00MB
tier:tier_nearline
tier_capacity:0.00MB

last_access_time 140604171325
parent_mdisk_grp_id:5
parent_mdisk_grp_name:p5
throttle_id:1
throttle_name:lcyfoxes_1
IOPs_limit:25000
bandwidth_limit_MB:500
volume_group_id:1
volume_group_name:ZlaIbra2
cloud_backup_enabled:no
cloud_account_id:
cloud_account_name:
backup_status:off
last_backup_time:
restore_status:available
backup_grain_size:

```

An invocation example

```
lsvdisk -delim : vv2
```

The resulting output:

```
id:0
name:vv2
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
formatting:yes
mdisk_id:many
mdisk_name:many
FC_id:
FC_name:
RC_id:
RC_name:
vdisk_UID:00000000000000AB:6005076801CF003F2800000000000000
```

```
preferred_node_id:1
fast_write_state:empty
cache:readwrite
udid:1234
fc_map_count:0
sync_rate:25
copy_count:2
se_copy_count:2
filesystem:
mirror_write_priority:latency
RC_change:no
compressed_copy_count:0
parent_mdisk_grp_id:5
parent_mdisk_grp_name:p5
encrypt:yes
volume_id:0
volume_name:vv2
function:master
```

```
copy_id:0
status:online
sync:yes
auto_delete:yes
primary:yes
mdisk_grp_id: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 tier0_flash
tier_capacity:1.63TB
```

```
tier:tier1_flash
tier_capacity:1.63TB
tier:tier_enterprise
tier_capacity:0.00MB
tier:tier_nearline
tier_capacity:0.00MB
compressed_copy:no
uncompressed_used_capacity:2.00GB
parent_mdisk_grp_id:5
parent_mdisk_grp_name:p5
```

```
copy_id:1
status:offline
sync:no
primary:no
mdisk_grp_id: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:off
easy_tier_status:inactive
```


```
tier tier0_flash
tier_capacity:1.63TB
tier:tier1_flash
tier_capacity:1.63TB
tier:tier_enterprise
tier_capacity:0.00MB
tier:tier_nearline
tier_capacity:0.00MB
compressed_copy:no
uncompressed_used_capacity:2.00GB
```

```
last_access_time 140604171325
parent_mdisk_grp_id:5
parent_mdisk_grp_name:p5
throttle_id:1
throttle_name:lcyfoxes_1
IOPs_limit:25000
bandwidth_limit_MB:500
volume_group_id:1
volume_group_name:ZlaIbra2
cloud_backup_enabled:no
cloud_account_id:
cloud_account_name:
backup_status:off
last_backup_time:
restore_status:available
backup_grain_size:
```

lsvdiskaccess

Use the **lsvdiskaccess** command to display a list of all input/output (I/O) groups in the volume access set.

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 displayed.

Note: Some filters allow the use of a wildcard when you enter the command. The following rules apply to the use of wildcards:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):
lsvdiskaccess -filtervalue "IO_group_name=io*"

-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?

(Optional) Displays the valid filter attributes for the **-filtervalue** parameter:

- IO_group_id
- IO_group_name
- scope
- status

vdisk_id | vdisk_name

(Optional) Specifies the volume for which to list access I/O groups.

Description

The **lsvdiskaccess** command lists the I/O groups in a volume access set. An accessible volume in an I/O group does not indicate the volume is mapped to any hosts. There is a detailed view and concise view, but the detailed view does not contain more information than the concise view.

This command returns values for the following volume attributes:

VDisk_id
Identifies the volume ID.

VDisk_name
Identifies the volume name.

IO_group_id
Identifies an I/O group ID in the volume access set.

IO_group_name
Identifies an I/O group name in the volume access set.

A detailed invocation example

```
lsvdiskaccess 0
```

The resulting output:

vdisk_id	vdisk_name	IO_group_id	IO_group_name
0	vdisk0	0	io_grp0
0	vdisk0	1	io_grp1
0	vdisk0	2	io_grp2

A concise invocation example

```
lsvdiskaccess
```

The resulting output:

vdisk_id	vdisk_name	IO_group_id	IO_group_name
0	vdisk0	0	io_grp0
0	vdisk0	1	io_grp1
0	vdisk0	2	io_grp2
3	vdisk3	1	io_grp1
7	vdisk7	0	io_grp0
7	vdisk7	2	io_grp2

lsvdiskanalysis

Use the **lsvdiskanalysis** command to display information for thin provisioning and compression estimation analysis report for a single volume or multiple volumes.

Syntax

```
➤➤ lsvdiskanalysis — [ -nohdr ] [ -delim delimiter ] —————➤➤
➤ [ -filtervalue — attribute=value ] [ -filtervalue? ] [ 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 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 1-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.

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 (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""), as follows:
`lsvdiskanalysis -filtervalue "usergrp_name=md*"`

-filtervalue?

(Optional) Displays the valid filter attributes for the **-filtervalueattribute=value** parameter:

- id
- name
- state

vdisk_id | vdisk_name

(Optional) Specifies the volume (by ID or name) to display compression estimation status for.

Description

This command displays information for thin provisioning and compression estimation analysis report for a single volume or multiple volumes.

Table 109 provides the attribute values that can be displayed as output view data.

Table 109. lsvdiskanalysis output

Attribute	Description
id	Indicates the ID (by number) of the object.
analysis_state	Indicates one of the following values: <ul style="list-style-type: none"> • idle, which means that the volume was never analyzed • scheduled, which means that the volume is scheduled for analysis (the analysis starts based on ascending volume IDs) • active, which means that the volume is being analyzed • estimated, which means that the volume was analyzed and the analysis results reflect the estimated savings from thin provisioning and compression • sparse, which means that the volume was analyzed but not enough samples of nonzero data were found • cancelling, which means the analysis is undergoing but there is a request to cancel it, and the analysis has not ended.
started_time	Indicates the date and time that the analysis started, which helps determine how long an estimate takes. The value must be an alphanumeric data string or be empty if analysis never started.

Table 109. *lsvdiskanalysis* output (continued)

Attribute	Description
analysis_time	Indicates the date and time that the analysis ended, to help determine how current the results are. If you cancel analysis, the value for the time is invalid (0, and it is not displayed). While active, time is invalid as well and does not reflect expected completion time. The value must be an alphanumeric data string or be empty if analysis never occurred.
capacity	Indicates the virtual capacity (host size) of the volume.
thin_size	Indicates the estimated size of the data without zero portions (thin-provisioned size).
thin_savings	Indicates how much data that is expected to be saved if it is a thin-provisioned volume.
thin_savings_ratio	Indicates the percent of data that is saved by thin-provisioned. The number must be a percentage. The number must be a percentage.
compressed_size	Indicates the estimated size of any nonzero data after compression completes.
compression_savings	Indicates how much data to expect to save if the volume is a compressed volume.
compression_savings_ratio	Indicates the amount of data that is saved by compression. The number must be a percentage.
total_savings	Indicates how much data to expect to save by converting a volume to a compressed volume.
total_savings_ratio	Indicates the amount of data saved by compression based on the overall volume capacity, which includes the thin nature of compressed volumes. The number must be a percentage.
accuracy	Indicates the accuracy estimation. The number must be a percentage.

A concise invocation example

```
lsvdiskanalysis
```

The detailed resulting output:

```
id name state analysis_time capacity thin_size thin_savings thin_savings_ratio compressed_size compression_savings compression_savings_ratio total_savings total_savings_ratio
0 ben0 idle 1.00GB 0.00MB 0.00MB 0 0.00MB 0.00MB 0 0.00MB 0
1 ben1 idle 1.00GB 0.00MB 0.00MB 0 0.00MB 0.00MB 0 0.00MB 0
2 ben2 active 1.00GB 0.00MB 0.00MB 0 0.00MB 0.00MB 0 0.00MB 0
3 ben3 idle 1.00GB 0.00MB 0.00MB 0 0.00MB 0.00MB 0 0.00MB 0
4 ben4 idle 1.00GB 0.00MB 0.00MB 0 0.00MB 0.00MB 0 0.00MB 0
5 ben5 idle 1.00GB 0.00MB 0.00MB 0 0.00MB 0.00MB 0 0.00MB 0
6 ben6 estimated 150608135456 1.00GB 62.18MB 961.82MB 93.92 12.23MB 49.95MB 80.33 1011.77MB 98.80
7 ben7 scheduled 1.00GB 0.00MB 0.00MB 0 0.00MB 0.00MB 0 0.00MB 0
8 ben8 idle 1.00GB 0.00MB 0.00MB 0 0.00MB 0.00MB 0 0.00MB 0
```

A detailed invocation example

```
lsvdiskanalysis
```

The detailed resulting output:

```
id 6
name ben6
state estimated
started_time 150608135446
analysis_time 150608135456
capacity 1.00GB
thin_size 62.18MB
thin_savings 961.82MB
thin_savings_ratio 93.92
compressed_size 12.23MB
compression_savings 49.95MB
compression_savings_ratio 80.33
total_savings 1011.77MB
total_savings_ratio 98.80
accuracy 4.97
```

lsvdiskanalysisprogress

Use the **lsvdiskanalysisprogress** command to display information about the space analysis progress for an entire clustered system (system).

Syntax

```
➤— lsvdiskanalysisprogress — [ -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 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 1-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 space analysis progress for an entire system.

Table 110 provides the attribute values that can be displayed as output view data.

Table 110. lsvdiskanalysisprogress output

Attribute	Description
vdisk_count	Indicates the number of volumes on this system.
pending_analysis	Indicates the number of volumes that belong to this system and: <ul style="list-style-type: none">• Scheduled a free space analysis• Have an active free space analysis• Are canceling free space analysis
estimated_completion_time	Indicates the estimated time at which analysis is expected to end. It is calculated based on number of scheduled volumes that are multiplied by 1 minute (there is no extrapolation from actual analysis duration). Estimated completion time does not consider volumes that are offline and displays estimated completion time as if the volumes are online.

A concise invocation example that shows progress for a system with some scheduled disks

```
lsvdiskanalysisprogress
```

The detailed resulting output:

```
vdisk_count pending_analysis estimated_completion_time
15          10              20150523135200
```

A concise invocation example that shows progress for a system with no scheduled disks

```
lsvdiskanalysisprogress
```


The detailed resulting output:

```
vdisk_count pending_analysis estimated_completion_time
15          0
```

lsvdiskcopy

Use the **lsvdiskcopy** command to list volume copy information.

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.

-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 volume 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 **lsvdiskcopy** 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 volume 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 volume are listed.

Description

The **lsvdiskcopy** command lists information for volume copies. If you specify the command with no parameters, all volumes and copies in the clustered system are listed.

The command returns values for the following volume copy attributes:

copy_id

Specifies a system-assigned identifier for the volume copy. The value can be 0 or 1.

status Indicates the system status. The value can be *online* or *offline*. A copy is offline if all nodes cannot access the storage pool that contains the copy.

sync Indicates whether the volume copy is synchronized.

auto_delete

Indicates the primary copy is deleted once the secondary copy is synchronized. The values are *yes* or *no*.

primary

Indicates whether the volume copy is the primary copy. A volume has exactly one primary copy. The value can be *yes* or *no*.

mdiskgrp_id

Indicates the ID of the storage pool that the volume copy belongs to.

mdiskgrp_name

Indicates the name of the storage pool that the volume copy belongs to.

type Indicates the virtualization type of the volume. The value can be *striped*, *sequential* or *image*.

mdisk_id

Indicates the MDisk ID that is used for sequential and image mode volumes.

mdisk_name

Indicates the MDisk name that is used for sequential and image mode volumes.

fast_write_state

Indicates the cache state of the volume copy. The value can be *empty*, *not_empty*, *corrupt*, or *repairing*. The value is always *empty* for non-thin-provisioned copies. A cache state of *corrupt* indicates that the volume is thin-provisioned and requires repair that is initiated by a **recovervdisk** command or the **repairsevdiskcopy** command.

used_capacity

Indicates the portion of *real_capacity* that is being used to store data. For non-thin-provisioned copies, this value is the same as the volume capacity. If the volume copy is thin-provisioned, the value increases from zero to the *real_capacity* value as more of the volume is written to.

Remember: This value is the same as the volume capacity value for fully-allocated copies.

real_capacity

Indicates the amount of physical storage that is allocated from an storage pool to this volume copy. If the volume copy is not thin-provisioned, the value is the same as the volume capacity. If the volume copy is thin-provisioned, the value can be different.

Remember: This value is the same as the volume capacity value for fully-allocated copies.

free_capacity

Indicates the difference between the `real_capacity` and `used_capacity` values.

Remember: This value is zero for fully-allocated copies.

overalllocation

Expressed as a percentage, indicates the ratio of volume capacity to `real_capacity` values. This value is always 100 for non-thin-provisioned volumes.

autoexpand

Indicates whether autoexpand is enabled on a thin-provisioned volume . The value can be on or off.

warning

Expressed as a percentage of the volume capacity for thin-provisioned or compressed volume copies, indicates that a warning is generated when the ratio of `used_capacity` to volume capacity reaches the specified level.

grainsize

For thin-provisioned volume copies, indicates the grain size chosen for the volume copy when it was created.

se_copy

Specifies if the copy is thin-provisioned.

easy_tier

Indicates 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

Indicates which Easy Tier functions are active for the volume copy:

- `active` indicates a pool is being managed by Easy Tier to provide tier management performance-based pool balancing. For example, extents of this volume copy can be moved for performance (automatic data placement).
- `inactive` indicates that no Easy Tier function is active.
- `balanced` indicates a pool is being managed by Easy Tier to provide performance-based pool balancing (for example, extents can be moved).
- `measured` indicates that statistics are being gathered for this volume copy, but no extents are moved.
- `balanced` indicates a pool is being managed by Easy Tier to provide performance-based pool balancing (for example, extents can be moved).

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 710)
Off	One	On	inactive (see note 2 on page 710)

Storage pool Easy Tier setting	Number of tiers in the storage pool	Volume copy Easy Tier setting	Volume copy Easy Tier status
Off	Two	Off	inactive
Off	Two	On	inactive (see note 2)
Measure	One	Off	measured (see note 3)
Measure	One	On	measured (see note 3)
Measure	Two	Off	measured (see note 3)
Measure	Two	On	measured (see note 3)
Auto	One	Off	measured (see note 3)
Auto	One	On	balanced (see note 4)
Auto	Two	Off	measured (see note 3)
Auto	Two	On	active (see note 5)
On	One	Off	measured (see note 3)
On	One	On	balanced (see note 4)
On	Two	Off	measured (see note 3)
On	Two	On	active (see note 5)

Note:

1. If the volume copy is in image or sequential mode or is being migrated then the volume copy Easy Tier status is 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 balanced, the Easy Tier function enables performance-based pool balancing for that volume copy.
5. When the volume copy status is active, the Easy Tier function operates in automatic data placement mode for that volume.
6. 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 except pool performance balancing are disabled for storage pools with a single tier, and that automatic data placement mode are enabled for all striped volume copies in a storage pool with two or more tiers.

tier Indicates which tier information is being reported:

- tier0_flash
- tier1_flash
- tier_enterprise
- tier_nearline

tier_capacity

Indicates the total MDisk capacity assigned to the volume in the tier.

Note: For thin-provisioned copies, the capacity by tier will be the real capacity.

compressed_copy

Indicates whether or not the volume copy is compressed.

uncompressed_used_capacity

For compressed volumes, indicates the amount of data written to the volume before compression.

parent_mdisk_grp_id

Indicates the physical storage pool id that the volume extents are actually allocated from. This is a numeric string (between 0 and 127 characters) or blank.

parent_mdisk_grp_name

Indicates the physical storage pool name that the volume extents are actually allocated from. This is an alphanumeric string (between 1 and 63 characters) or blank.

encrypt

Indicates whether the volume and its copies are encrypted. The values are yes or no.

An invocation example

```
lsvdiskcopy -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:compressed_copy
0:RAM_V2:0:online:yes:yes:2:RAM_MDG2:5.00GB:striped:yes:on:inactive:yes
1:RAM_V3:0:online:yes:yes:2:RAM_MDG2:5.00GB:striped:no:on:inactive:no
2:RAM_V4:0:online:yes:yes:1:RAM_MDG3:5.00GB:striped:no:on:inactive:yes
3:RAM_V5:0:online:yes:yes:2:RAM_MDG2:5.00GB:striped:yes:on:inactive:no
3:RAM_V5:1:online:yes:no:2:RAM_MDG2:5.00GB:striped:yes:on:inactive:yes
4:RAM_V1:0:online:yes:yes:3:RAM_MDG1:5.00GB:striped:no:on:inactive:no
5:RAM_V6:0:online:yes:yes:0:RAM_MDG4:5.00GB:striped:yes:on:inactive:yes
```

An invocation example

```
lsvdiskcopy -copy 0 -delim : vv1
```

The resulting output:

```
vdisk_id:0
vdisk_name:vv1
capacity:16.00GB
copy_id:0
status:online
sync:yes
auto_delete: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:tier0_flash  
tier_capacity:1.63TB  
tier:tier1_flash  
tier_capacity:1.63TB  
tier:tier_enterprise  
tier_capacity:0.00MB  
tier:tier_nearline  
tier_capacity:0.00MB  
tier_capacity:64.00MB
```

```
tier:tier0_flash  
tier_capacity:1.63TB  
tier:tier1_flash  
tier_capacity:1.63TB  
tier:tier_enterprise  
tier_capacity:0.00MB  
tier:tier_nearline  
tier_capacity:0.00MB  
tier_capacity:7.94GB  
compressed_copy:yes  
uncompressed_used_capacity:1.0MB  
parent_mdisk_grp_id:5  
parent_mdisk_grp_name:p5  
encrypt:yes
```

An invocation example

```
lsvdiskcopy -copy 0 -delim : vv1
```

The resulting output:

```
vdisk_id:0  
vdisk_name:vv1  
capacity:16.00GB  
copy_id:0  
status:online  
sync:yes  
auto_delete:yes  
primary:yes  
mdisk_grp_id: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:tier0_flash  
tier_capacity:1.63TB  
tier:tier1_flash  
tier_capacity:1.63TB  
tier:tier_enterprise  
tier_capacity:0.00MB  
tier:tier_nearline  
tier_capacity:0.00MB  
compressed_copy:no
```

```
uncompressed_used_capacity:8.00GB
parent_mdisk_grp_id:5
parent_mdisk_grp_name:p5
encrypt:yes
```

lsvdiskdependentmaps

Use the **lsvdiskdependentmaps** command to display all FlashCopy mappings with target volumes that are dependent upon data held on the specified volume.

Syntax

```
➤— lsvdiskdependentmaps — [ -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 name or ID of a volume.

Description

The **lsvdiskdependentmaps** command displays FlashCopy mappings that have target volumes 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 volume *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

```
lsvdiskdependentmaps -delim : 0
```


The concise resulting output:

```
id:name
2:fcmap2
5:fcmap5
```

lsvdiskextent

Use the **lsvdiskextent** command to list the MDisk extents provided for the specified volumes.

Syntax



Parameters

-copy *copy_id*

(Optional) Displays a list of MDisk IDs that are members of the specified volume 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 volume 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 volumes.

Each volume is constructed from one or more MDisk. To determine the relationship between a volume and its MDisk, issue the following command:

```
lsvdiskmember vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the volume. This command displays a list of MDisk IDs that make up the volume.

To determine the number of extents that are provided by each MDisk, issue the following command:

```
lsvdiskextent vdisk_name | vdisk_id
```

where *vdisk_name* | *vdisk_id* is the name or ID of the volume. This command displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the given volume.

To determine the relationship between MDisk and volumes, issue the following command for each MDisk:

```
lsmdiskmember 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 volumes that are using this MDisk.

To determine the relationship between MDisks and volumes, and the number of extents that are used by each volume, you must use the command-line interface. For each MDisk, issue the following command:

```
lsmdiskextent mdisk_name | mdisk_id
```

where *mdisk_name* | *mdisk_id* is the name or ID of the MDisk. This command displays a table of volume IDs and the corresponding number of extents that are used by each volume.

An invocation example

```
lsdiskextent -delim : vdisk0
```

The resulting output

```
id:number_extents
0:0
```

lsdiskfcmappcopies

Use the **lsdiskfcmappcopies** command to display a list of all FlashCopy mappings with a target volume containing a valid copy of the specified volume.

Syntax

```
➡➡ lsdiskfcmappcopies — [ -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 volume for which the FlashCopy mappings are displayed.

Description

This command returns a list of the FlashCopy mappings that have a target volume with a valid copy of the specified volume. The target volumes of these mappings can be considered as candidate source volumes for mappings to restore from.

The mappings returned are in the copying, idle_copied, or stopping state with 100% progress.

Note: Maps that are rc_controlled are not shown in the view when this command is specified.

An invocation example

```
lsdiskfcmappings -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
```

lsdiskfcmappings

Use the **lsdiskfcmappings** command to display a list of FlashCopy mappings to which the volume belongs. A volume can be part of up to 256 FlashCopy mappings.

Syntax

```
►► lsdiskfcmappings — [ -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 volume for which a list of all FlashCopy mappings is required.

Description

The **lsdiskfcmappings** command returns a list of all FlashCopy mappings that the volume is a member of. The list is returned in no particular order.

An invocation example

```
lsdiskfcmappings -delim : vdisk2
```

The resulting output:

```
fc_id:fc_name
1:fcmap1
3:fcmap3
```

lsvdiskhostmap

Use the **lsvdiskhostmap** command to list the volumes to the host mapping. These hosts have the specified volumes mapped to them; the volume is visible to these hosts.

Syntax

```
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 volume. The clustered system displays a list of all the hosts to which this volume is mapped and the Small Computer System Interface (SCSI) ID by which the volume is mapped.

Description

This command displays a list of host IDs and names. These hosts have the specified volume mapped to them; that is, the volume is visible to these hosts. The SCSI LUN ID is also displayed. The SCSI LUN ID is the ID by which the volume is recognized by the host.

Determining the host that a volume is mapped to: List the hosts that this volume is mapped to, by issuing the following command:

```
lsvdiskhostmap vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the volume. A list is displayed. Look for the host name or ID to determine which host this volume is mapped to. If no data is displayed, the volume is not mapped to any hosts.

The command returns the following values:

id Specifies the ID of the volume in the output for **lsvdiskhostmap**.

name Specifies the name of the volume in the output for **lsvdiskhostmap**.

SCSI_id

Specifies the SCSI ID.

host_id
Specifies the ID of the host.

host_name
Specifies the name of the host.

vdisk_UID
Specifies the UID of the volume.

IO_group_id
Specifies the ID of the input/output (I/O) group in which the host volume mapping exists.

IO_group_name
Specifies the name of I/O group in which the host volume mapping exists.

host_cluster_id
Indicates the unique ID for a host system.

host_cluster_name
Indicates the unique name for a host system.

An invocation example

```
lsvdiskhostmap 0
```

The resulting output:

id	name	SCSI_id	host_id	host_name	vdisk_UID	IO_group_id	IO_group_name	hostcluster_id	hostcluster_name
4	vdisk4	0	3	host3	UID4	0	iogrp0		
6	priv_6	4	4	host4	UID6	0	iogrp0		
8	shared_8	5	4	host4	UID8	0	iogrp0	0	hostcluster0
8	shared_8	5	5	host5	UID8	0	iogrp0	0	hostcluster0

lsvdisklba

Use the **lsvdisklba** command to list the volume and logical block address (LBA) for the specified storage pool LBA.

Syntax

```

>>> lsvdisklba -- -mdisklba -- mdisklba -- [ -delim -- delimiter ] [ - nohdr ]
> -mdisk [ mdisk_id [ mdisk_name ]

```

Parameters

-mdisklba mdisklba

(Required) Specifies the 64-bit hexadecimal LBA on the MDisk. The LBA must be specified in hex, with a 0x prefix.

-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 volume that is associated with the MDisk LBA.

If applicable, the command also lists the range of LBAs on both the volume and MDisk that are mapped in the same extent, or for thin-provisioned disks, in the same grain.

The **vdisk_lba** field provides the corresponding LBA on the virtual capacity for the input LBA. For compressed volume copies it is blank, and the system provides the ranges of virtual LBAs that are compressed into the input LBA.

Table 111 provides command output that depends on several variables.

Table 111. *lsvdisklba* command output scenarios

Field	Typical scenario	Quorum disk	Thin-provisioned metadata	Extent not allocated	Formatting extent	Extent allocated to thin-provisioned disk, LBA not used on thin-provisioned disk
copy_id	yes	no	yes	no	yes	yes
vdisk_id	yes	no	yes	no	yes	yes
vdisk_name	yes	no	yes	no	yes	yes
type	allocated	metadata	metadata	unallocated	formatting	unallocated
vdisk_lba	yes	no	no	no	no	no
vdisk_start	yes	no	no	no	no	no
vdisk_end	yes	no	no	no	no	no
mdisk_start	yes	yes	yes	yes	yes	yes
mdisk_end	yes	yes	yes	yes	yes	yes

An invocation example

```
lsvdisklba -mdisk 1 -mdisklba 0x100123
```


The resulting output:

```
vdisk_id vdisk_name copy_id type      vdisk_lba  vdisk_start vdisk_end  mdisk_start      mdisk_end
0         vdisk0      0         allocated 0x00000123 0x00000000 0x000FFFFF 0x0000000000100000 0x00000000001FFFFF
```

lsvdiskmember

Use the **lsvdiskmember** command to display a list of MDisks that are members of the specified volume.

Syntax



Parameters

-copy *copy_id*

(Optional) Displays a list of MDisks that are members of the specified volume 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) Identifies the specific volume to query..

Description

This command displays a list of managed disks, which provide extents that make up the volume that is specified by the ID.

Every volume 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 **lsvdiskmember** command, the concise view displays a list of volumes. These are the volumes 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 volumes and MDisks, issue the following command:

```
lsvdiskmember vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the volume. This displays a list of IDs that correspond to the MDisks that make up the volume.

To determine the relationship between volumes and MDisks, and the number of extents that are provided by each MDisk, you must use the command-line interface. Issue the following command:

```
lsvdiskextent vdisk_id | vdisk_name
```

where *vdisk_id* | *vdisk_name* is the name or ID of the volume. This displays a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the specified volume.

To determine the relationship between MDisks and volumes, issue the following command:

```
lsmdiskmember 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 volumes that are using this MDisk.

To determine the relationship between MDisks and volumes, and the number of extents that are used by each volume, you must use the command-line interface. For a specified MDisk, issue the following command:

```
lsmdiskextent mdisk_id | mdisk_name
```

where *mdisk_id* | *mdisk_name* is the name or ID of the MDisk. This displays a table of volume IDs and the corresponding number of extents that are used by each volume.

An invocation example

```
lsdiskmember 1
```

The resulting output:

```
id
2
```

lsdiskprogress

Use the **lsdiskprogress** command to track the progress during new volume formatting.

Syntax

```
➤— lsdiskprogress — [ -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.

(Optional) Specifies the volume ID or name. If you do not specify this parameter, the progress of all volumes currently being formatted is displayed.

This command displays the progress of the format of a new volume as a completed percentage. If the volume has multiple copies, the command reports the average progress of the format.

id	Indicates the ID of the volume being formatted.
-----------	---

Indicates the formatting progress.

Indicates the estimated time in which the formatting operation completes. The value is in the YYYYMMDDHHMMSS format, and is blank if the duration is not known.

```
lsvdiskprogress -delim : 0
```

```
id:0
progress:58
estimated completion time:150101010100
```

Use the **lsvdisksyncprogress** command to display the progress of volume copy synchronization.

```

▶▶ lsvdisksyncprogress — [ -nohdr ] [ -delim delimiter ] [ -copy id ]
▶ [ -vdisk_name ] [ -vdisk_id ]

```

(Optional) Specifies the volume copy ID to list synchronization progress. You must also specify a *vdisk_name* | *vdisk_id* value. If you do not specify this parameter, progress is displayed for all copies.

(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.

(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*

(Optional) Specifies the volume name or ID to list synchronization progress.

Description

To display the volume copies that require synchronization, specify the command with no parameters. To display the synchronization progress for all copies of a volume, specify the command with the *vdisk_name* | *vdisk_id* parameter. Estimated completion time is displayed in the YYMMDDHHMMSS format. The command displays progress for the following special cases as:

- A synchronized copy displays a progress of 100 and a blank estimated completion time.
- An offline copy or a copy with a zero synchronization rate displays a blank estimated completion time. An offline copy displays (gradually) decreasing progress if the volume is being written to.
- Nonmirrored volumes 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 volume synchronization. After you create a mirrored volume using the **mkvdisk** or **addvdiskcopy** command, you can use the command to monitor the progress of the synchronization.

An invocation example

```
lsvdisksyncprogress
```

The resulting output

vdisk_id	vdisk_name	copy_id	progress	estimated_completion_time
0	vdisk0	1	50	070301150000
3	vdisk3	0	72	070301132225
4	vdisk4	0	22	070301160000
8	vdisk8	1	33	

An invocation example

```
lsvdisksyncprogress vdisk0
```

The resulting output

vdisk_id	vdisk_name	copy_id	progress	estimated_completion_time
0	vdisk0	0	100	
0	vdisk0	1	50	070301150000

Isvolumebackup

Use the **lsvolumebackup** command to list the volumes that have cloud snapshot that enabled and volumes that have cloud snapshots in the cloud account.

Syntax

```

>>> lsvolumebackup — [ -filtervalue — attribute_value ] [ -nohdr ]
> [ -delim — delimiter ] [ -filtervalue? ]

```

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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""): `lsvolumebackup -filtervalue volume_id="1*"`

-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?

(Optional) Displays a list of valid filter attributes. The following filters for the **lsvolumebackup** command are valid:

- `volume_uid`
- `volume_id`
- `volume_name`
- `volume_group_id`
- `volume_group_name`
- `cloud_account_id`
- `cloud_account_name`

Description

This command lists the volumes that have cloud snapshot enabled and also lists volumes that have cloud snapshots in the cloud account.

This view spans all cloud accounts. To refresh the view by reloading what is on the cloud, specify `chcloudaccountawss3 -refresh 0` or `chcloudaccountswift -refresh 0`.

A snapshot that is being copied to the cloud (which means the snapshot status value is `copying` or `copying_error`) does not count towards the generation count total for the volume. This means it is not available for a restore, and if there is a local system failure, the generation no longer exists in the cloud.

The specified volume appears in the displayed output when the volume has cloud snapshot that is enabled. The generation count is initially 0 and remains 0 while the volume copy is in progress. The generation count changes to 1 after the copy completes.

The last snapshot time is blank while the first snapshot is in progress. A snapshot that is being deleted from the cloud counts toward the generation count for the volume even if it is not available for a restore.

This table provides the attribute values that can be displayed as output view data.

Table 112. Isvolumebackup output

Attribute	Description
volume_UID	Indicates the volume UID.
volume_id	Indicates the volume ID if a volume with the specified UID exists on the local system. The value must be a number (or blank).
volume_name	Indicates the volume name. The value must be an alphanumeric string.
volume_group_id	Indicates the volume group ID that the volume is a member of (if applicable), if a volume with the specified UID exists on the local system. The value must be a number (or blank).
volume_group_name	Indicates the volume group name that the volume is a member of (if applicable), if a volume with the specified UID exists on the local system. The value must be an alphanumeric string (or blank).
cloud_account_id	Indicates the ID for the cloud account that contains the volume snapshots.
cloud_account_name	Indicates the cloud account name for the cloud account that contains the volume snapshot. The value must be an alphanumeric string.
last_backup_time	Indicates the timestamp of the most recent snapshot for this volume. The value must be in YYMMDDHHMMSS format or blank.
generation_count	Indicates the number of snapshot generations that exist for the specified volume. The value must be a number. Note: <ul style="list-style-type: none"> Any generations that are being copied to the cloud do not count towards this number. Any generations that are being deleted count towards this number until the delete process completes.
backup_size	Indicates the approximate amount of storage (the capacity) that is in use by snapshot generations for the specified volume

An invocation example

`lsvolumebackup`

The resulting output:

```

volume_UID      volume_id volume_name volume_group_id volume_group_name cloud_account_id cloud_account_name
600507680CA880DF18000000000000002 2      vdisk2      2      logArchive      0      myAmazon
600507680CA880DF18000000000000003 3      vdisk3      2      logArchive      0      myAmazon
600507680CA880DF18000000000000004 4      vdisk4      2      logArchive      0      myAmazon
600507680CA880DF18000000000000017 17     vdisk4      2      logArchive      0      myAmazon

```

lsvolumebackupgeneration


Use the `lsvolumebackupgeneration` command to list any volume snapshots available on the specified volume.

Syntax

```

➤— lsvolumebackupgeneration — [ -filtervalue — attribute_value ] — [ -nohdr ] — ➤

```



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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (").

-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 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 1-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 valid filter attributes. The following filters for the **lsvolumebackupgeneration** command are valid:

- **state**

-volume *volume_name* | *volume_id*

(Required) Specifies the volume to list cloud snapshots for by volume ID or name.

Note: The volume must exist on the local system.

The value for volume ID must be a number and the value for volume name must be an alphanumeric string. This parameter is mutually exclusive with **-uid**.

-uid *volume_UID*

(Optional) Specifies the volume to list cloud snapshots for by volume UID. This parameter is mutually exclusive with **-volume**.

Description

This command lists any available volume snapshots for the specified volume.

This table provides the attribute values that can be displayed as output view data.

Table 113. *lsvolumebackupgeneration* output

Attribute	Description
generation_id	Indicates the snapshot generation volume ID. The value must be a number.
backup_time	Indicates the timestamp of the most recent snapshot. The value must be in YYMMDDHHMMSS format (or blank).
volume_group_name	Indicate the volume group name. The value must be an alphanumeric string (or blank).
volume_size	Indicates the virtual capacity of the volume during snapshot generation. This value can differ from the current volume size (the capacity in MB or GB) if a volume with the specified UID exists on the local system.
type	Indicates the type of volume snapshot generation. The values are full and incremental.
state	Indicates the state of the volume backup generation (in the cloud system). The values are: <ul style="list-style-type: none"> • copying • complete • deleting
cloud_upload_size	Indicates the amount of data (the capacity in MB or GB) that is uploaded from the snapshot generation volume to the cloud system.

Generations that are being copied to the cloud account are included in the copying state value. Generations that are being deleted from the cloud account are also included in the deleting state value.

An invocation example

```
lsvolumebackupgeneration -volume 2
```

The resulting output:

generation_id	backup_time	volume_group_name	volume_size	type	state	cloud_upload_size
1	160217021250		50.00GB	full	complete	2.83GB
2	160217021355		50.00GB	incremental	complete	177.50MB
3	160218021402		50.00GB	incremental	complete	132.02MB
4	160219021400		50.00GB	incremental	copying	12.43MB

An invocation example

```
lsvolumebackupgeneration -uid 600507680CA880AB12000000000000015
```

The resulting output:

generation_id	backup_time	volume_group_name	volume_size	type	state	cloud_upload_size
1	160215021355		10.00GB	full	complete	53.88MB

lsvolumebackupprogress

Use the **lsvolumebackupprogress** command to display information about the progress of snapshot operations.

Syntax

```

>>> lsvolumebackupprogress -- [ -filtervalue -- attribute_value ] [ -nohdr ]
[ -delim -- delimiter ] [ -filtervalue? ]

```

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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):
`lsvolumebackupprogress -filtervalue volume_id="1*"`

-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?

(Optional) Displays a list of valid filter attributes. The following filters for the **lsvolumebackupprogress** command are valid:

- volume_UID
- volume_id
- volume_name
- task
- status

Description

This command display information about the progress of snapshot operations.

This table provides the attribute values that can be displayed as output view data.

Table 114. lsvolumebackupprogress output

Attribute	Description
volume_UID	Indicates the volume UID. The value must be a number 0 - 32 characters long.
volume_id	Indicates the volume ID (if a volume with the specified UID exists on the local system). The value must be a number (or blank).
volume_name	Indicates the volume name (if a volume with the specified UID exists on the local system). The value must be a an alphanumeric string (or blank).
task	Indicates the type of task that is in progress. The values are backup and delete.

Table 114. *lsvolumebackupprogress* output (continued)

Attribute	Description
status	Indicates the task status. The values are: <ul style="list-style-type: none"> • copying • copying_error • deleting • deleting_error
generation_id	Indicates the generation ID for the volume being created or deleted. The value must be a number.
backup_time	Indicates the snapshot time for the volume being copied to the cloud system. The value must be in YYMMDDHHMMSS format for snapshot tasks or blank for deletion tasks.
progress	Indicates the task progress as a percentage. The value must be a number 0 - 99.
error_sequence_number	Indicates a particular error number. The value must be a number (or blank).

An invocation example

`lsvolumebackupprogress`

The resulting output:

volume_UID	volume_id	volume_name	task	status	generation_id	backup_time	progress	error_s
600507680CA880DF1800000000000002	2	vdisk2	backup	copying	6	160218191005	88	
600507680CA880DF18000000000000015	15	vdisk15	backup	copying_error	19	160218190845	12	122
600507680CA880DF18000000000000015	15	vdisk15	delete	deleting	8		5	
600507680CA880DF18000000000000017		vdisk108	delete	deleting	10		17	
600507680CA880DF18000000000000018		vdisk109	delete	deleting	many		55	

lsvolume group

Use the `lsvolume group` command to display information about configured volume groups.

Syntax

```

▶▶ lsvolume group — [ -filtervalue — attribute_value ] [ -nohdr ]
▶ [ -delim — delimiter ] [ -filtervalue? ] [ volumegroup_id — volumegroup_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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):

```
lsvolume group -filtervalue id="1*"
```

-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 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 1-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.

-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.

-filtervalue?

(Optional) Displays a list of valid filter attributes. The following filters for the **lsvolume group** command are valid:

- id
- name
- volume_count
- backup_status
- last_backup_time

volume group_id | volume group_name

(Optional) Specifies a volume group ID or volume group name. The value must be a number for the ID and an alphanumeric string for the name.

Description

This command displays information about configured volume groups.

This table provides the attribute values that can be displayed as output view data.

Table 115. Isvolume group output

Attribute	Description
id	Indicates the volume group ID. The value must be a number.
name	Indicates the volume group name. The value must be an alphanumeric string.
volume_count	Indicates the number of volume members in a group. The value must be a number.
backup_status	Indicates whether a new (volume) group snapshot can be started. The values are: <ul style="list-style-type: none"> • empty • off • not_ready • ready • copying • copying_error If a group snapshot is in progress, the value represents the snapshot operation status.
last_backup_time	Indicates the most recent volume group snapshot time. The value must be in the YYMMDDHHMMSS format (or blank).

A concise invocation example

```
lsvolumegroup
```

The resulting output:

id	name	volume_count	backup_status	last_backup_time
0	VG1	5	copying	160308115216
1	VG2	0	not_ready	150408115216

A detailed invocation example

```
lsvolumegroup 1
```

The resulting output:

```
id 1
name VG2
volume_count 0
backup_status not_ready
last_backup_time
```

Isvolumerestoreprogress

Use the **lsvolumerestoreprogress** command to display information about restore operation progress.

Syntax

```
➤— lsvolumerestoreprogress — [ -filtervalue — attribute_value ] [ -nohdr ]
[ -delim — delimiter ] [ -filtervalue? ] [ volume_name ]
[ volume_id ] ➤
```

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:

- The wildcard character is the asterisk (*).
- The command can contain a maximum of one wildcard.
- When you use a wildcard, enclose the filter entry within double quotation marks (""):

```
lsvolumerestoreprogress -filtervalue volume_id="1*"
```

-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?

(Optional) Displays a list of valid filter attributes. The following filters for the **lsvolumerestoreprogress** command are valid:

- volume_id
- volume_name
- task
- status

`volume_name | volume_id`

(Optional) Indicates the volume name or ID for the volume being restored. The value for volume name must be an alphanumeric string and the value for volume ID must be a number.

Description

This command displays information about restore operation progress.

This table provides the attribute values that can be displayed as output view data.

Table 116. lsvolumerestoreprogress output

Attribute	Description
volume_id	Indicates the volume ID for the volume being restored. The value must be a number (or blank).
volume_name	Indicates the volume name for the volume being restored. The value must be an alphanumeric string (or blank).
task	Indicates the type of task that is in progress. The value is restore.
status	Indicates the task status. The values are: <ul style="list-style-type: none"> • restoring • restoring_error
generation_id	Indicates the generation ID for the volume snapshot generation being restored. The value must be a number.
backup_time	Indicates the time for the volume snapshot generation that is being restored to the cloud system (or available on the restore volume). The value must be in YYMMDDHHMMSS format for snapshot tasks or blank for deletion tasks.
progress	Indicates the task progress as a percentage. The value must be a number 0 - 100. The status is blank when the status is uncommitted.
error_sequence_number	Indicates a particular error number. The value must be a number (or blank).
volume_backup_id	Indicates the UID of the volume snapshot being restored. The value must be a number 1 - 32.
restore_volume_id	Indicates the ID of the volume that is the target of the restore operation. This volume is either the production volume (which means that the restore volume ID is the same as the volume ID) or a temporarily restored volume that is automatically provisioned by the restore process. The value must be a number.
restore_volume_name	Indicates the name of the volume that is the target of the restore operation. This volume is either the production volume (which means that the restore volume name is the same as the volume name) or a temporarily restored volume that is automatically provisioned by the restore process. The value must be an alphanumeric string.

A concise invocation example

```
lsvolumerestoreprogress
```

The resulting output:

volume_id	volume_name	task	status	generation_id	backup_time	progress	error_sequence_number
2	vdisk2	restore	restoring	17	160102104511	88	
21	vdisk21	restore	restoring_error	4	160102105023	19	

122

A detailed invocation example

```
lsvolumerestoreprogress vdisk2
```

The resulting output:

```
volume_id 2
volume_name vdisk2
task restore
status restoring
generation_id 17
backup_time 160102104511
progress 88
error_sequence_number
volume_backup_UID 600507680CA880DF18000000000000002
restore_volume_id 2
restore_volume_name vdisk2
```

mkmetadatavdisk

Use the **mkmetadatavdisk** command to create one metadata volume (owner type is `host_integration_metadata`) from a storage pool. You can also export one block device or file system (that is based on this volume) in the configuration node.

Syntax

```
➤— mkmetadatavdisk — — -mdiskgrp ———— mdiskgrp_id —————➤
                                   |
                                   ———— mdiskgrp_name ————
```

Parameters

-mdiskgrp *mdiskgrp_id* | *mdiskgrp_name*

(Required) Assigns one or multiple storage pools for use in creating a metadata volume. The value must be a numeric value for *mdiskgrp_id* and an alphanumeric string for *mdiskgrp_name*.

Description

This command creates one metadata volume from a storage pool.

An invocation example

```
mkmetadatavdisk -mdiskgrp pool_a
```

The resulting output:

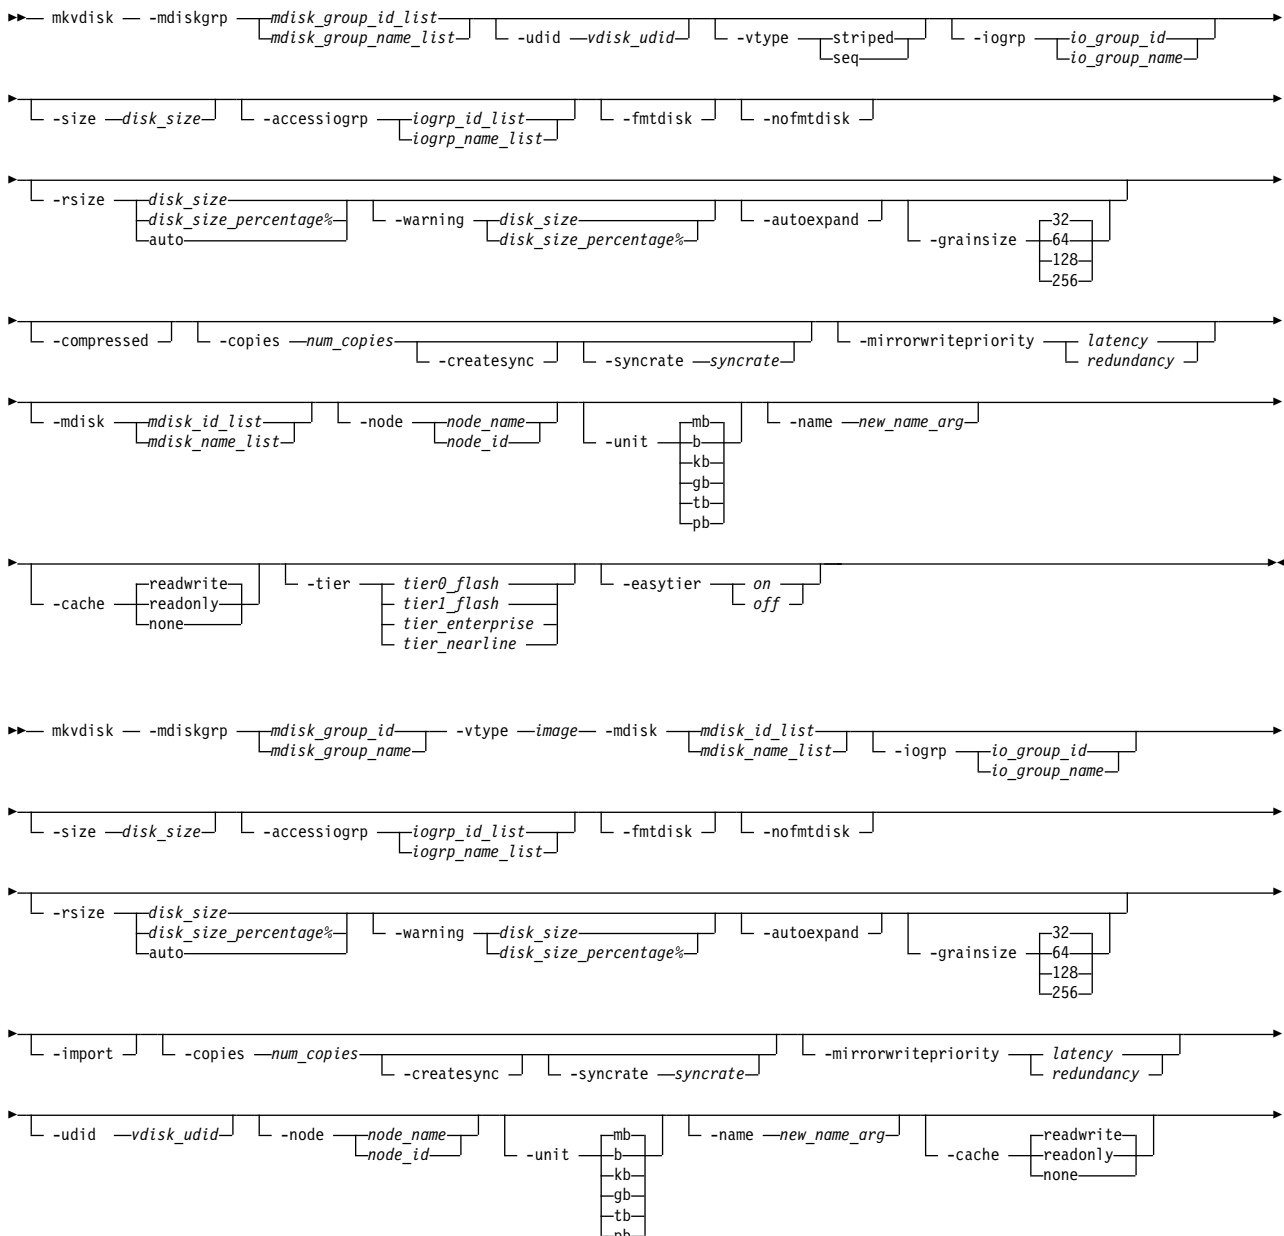
```
Virtual Disk, id [2], successfully created
```


mkvdisk

Use the **mkvdisk** command to create sequential, striped, or image mode volume objects. When they are mapped to a host object, these objects are seen as disk drives with which the host can run I/O operations. Note that the first syntax diagrams below is for striped or sequential volumes and the second syntax diagram is for image mode volumes. Use the **mkvolume** command for a simplified way of creating high availability volumes. This includes stretched and hyperswap topologies. Use the **mkimagevolume** command for a simplified way of creating an image mode volume, importing existing data from a managed disk.

Note: The first syntax diagram depicts the creation of a **sequential** or **striped** mode volume. The second syntax diagram depicts the creation of an **image** mode volume.

Syntax





Parameters

-mdiskgrp *mdisk_group_id_list* | *mdisk_group_name_list*

(Required) Specifies one or more storage pools to use when you are creating this volume. If you are creating multiple copies, you must specify one storage pool per copy. The primary copy is allocated from the first storage pool in the list.

-iogrp *io_group_id* | *io_group_name*

(Optional) Specifies the I/O group (node pair) with which to associate this volume. If you specify **-node**, you must also specify **-iogrp**.

Remember:

- Create the first compressed volume copy for an I/O group to activate compression.
- You cannot create or move a volume copy that is compressed to an I/O group that contains at least one node that does not support compressed volumes. You must select another I/O group to move the volume copy to (but this does not affect moving to the recovery I/O group).

-accessiogrp *iogroup_id_list* | *iogroup_name_list*

(Optional) Specifies the members of the volume I/O group access set. If this option is not specified, only the caching I/O group is added to the volume I/O group access set. If any access I/O groups are specified, only those I/O groups are in the access set (including if that set does not include the caching I/O group).

-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 or striped volume creation) (Optional for image volume creation) Specifies the capacity of the volume, 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. This 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 thin-provisioned or compressed volume. This parameter makes the volume thin-provisioned; otherwise, the volume is fully-allocated. Specify the *disk_size* | *disk_size_percentage* value by using an integer, or an integer immediately followed by the percent character (%). Specify the units for a *disk_size* integer by using the **-unit** parameter; the default is MB. The **-rsize** value can be greater than, equal to, or less than the size of the volume. The **auto** option creates a volume copy that uses the entire size of the MDisk; if you specify the **-rsize auto** option, you must also specify the **-vtype image** option.

-fmtdisk

(Optional) Specifies that the volume be formatted. This parameter is no longer required for any volumes.

This parameter is not required when creating fully-allocated volumes. The format operation is automatically applied to fully-allocated volumes unless you specify **-nofmtdisk** parameter. The format

operation sets the extents that make up this volume to all zeros after it is created. This process takes place in the background concurrently with host I/O operations on the new volume.

Remember: Formatting is on by default for single copy, fully-allocated, and non-image mode volumes. You cannot format an image mode volume.

The format operation completes asynchronously. You can query the status by using the **lsvdiskprogress** command. You cannot specify this parameter with the **-vtype image** parameter.

This parameter is not required when you create thin-provisioned volumes. Thin-provisioned volumes return zeros for extents that are not written to. No format operation is required. This parameter also synchronizes mirrored copies by default.

-nofmtdisk

(Optional) Specifies that formatting be turned off for the new volume.

Remember: Formatting is on by default for single copy, fully-allocated, and non-image mode volumes, and you can specify this parameter to turn it off.

-compressed

(Optional) Enables compression for the volume. This parameter must be specified with **-rsize** and cannot be specified with **-grainsize**.

-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 volume copies. A warning is generated when the used disk capacity on the thin-provisioned 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 volume size.

Important: If **-autoexpand** is:

1. Enabled, the default value for **-warning** is 80% of the volume capacity.
2. Not enabled, the default value for **-warning** is 80% of the real capacity.

To disable warnings, specify **0**.

-autoexpand

(Optional) Specifies that thin-provisioned copies automatically expand their real capacities by allocating new extents from their storage pool. 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 storage pool runs out of space by having the storage pool to consume this reserved space first.

The parameter has no immediate effect on image mode copies. However, if the image mode copy is later migrated to managed mode, the copy is then automatically expanded.

-grainsize 32 | 64 | 128 | 256

(Optional) Sets the grain size (KB) for a thin-provisioned volume. This parameter also requires that the **-rsize** parameter be specified. If you are using the thin-provisioned volume in a FlashCopy map, use the same grain size as the map grain size for best performance. If you are using the thin-provisioned volume directly with a host system, use a small grain size. The grain size value must be 32, 64, 128, or 256 KB. The default is 256 KB.

-import

(Optional) Imports a thin-provisioned volume 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 volume. The default value is 1.

-syncrate *syncrate*

(Optional) Specifies the copy synchronization rate. A value of zero (0) prevents synchronization. The default value is 50. See Table 117 on page 741 for the supported **-syncrate** values and their corresponding rates. Use this parameter to alter the rate at which the fully-allocated volume or mirrored volume format before synchronization.

-createsync

(Optional) Creates copies in sync. Use this parameter if you have already formatted the MDisk, or when read stability to unwritten areas of the volume is not required.

-mirrorwritepriority *latency* | *redundancy*

(Optional) Specifies how to configure the mirror write algorithm priority. If not specified, the default value is *latency*.

1. Choosing *latency* means a copy that is slow to respond to a write input/output (I/O) becomes unsynchronized, and the write I/O completes if the other copy successfully writes the data.
2. Choosing *redundancy* means a copy that is slow to respond to a write I/O synchronizes completion of the write I/O with the completion of the slower I/O to maintain synchronization.

-vtype *seq* | *striped* | *image*

(Optional) Specifies the virtualization type. When creating sequential or image mode volumes, 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 volume. You can use the **-node** parameter to specify the preferred access node. If you specify **-node**, you must also specify **-iogrp**.

Note: This parameter is required for the subsystem device driver (SDD). The system 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. The default unit type is MB.

-mdisk *mdisk_id_list* | *mdisk_name_list*

(Optional) Specifies one or more managed disks. For sequential and image mode volumes, the number of MDisk must match the number of copies. For sequential mode volumes, each MDisk must belong to the specified storage pool. For striped volumes, you cannot specify the **-mdisk** parameter if the **-copies** value is greater than 1.

When creating a single copy striped volume, you can specify a list of MDisk to stripe across.

You must use this parameter to specify an MDisk that has a mode of unmanaged.

-name *new_name_arg*

(Optional) Specifies a name to assign to the new volume.

-cache *readwrite* | *readonly* | *none*

(Optional) Specifies the caching options for the volume. Valid entries are:

- *readwrite* enables the cache for the volume
- *readonly* disables write caching while allowing read caching for a volume
- *none* disables the cache mode for the volume

The default is *readwrite*.

Remember: If you do not specify the **-cache** parameter, the default value (*readwrite*) is used.

-tier *tier0_flash* | *tier1_flash* | *tier_enterprise* | *tier_nearline*

(Optional) Specifies the MDisk tier when an image mode copy is added.

tier0_flash

Specifies a tier0_flash hard disk drive or an external MDisk for the newly discovered or external volume.

tier1_flash

Specifies an tier1_flash (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

tier_enterprise

Specifies a tier_enterprise hard disk drive or an external MDisk for the newly discovered or external volume.

tier_nearline

Specifies a tier_nearline hard disk drive or an external MDisk for the newly discovered or external volume.

ssd

Specifies an SSD (or flash drive) hard disk drive or an external MDisk for the newly discovered or external volume.

nearline

Specifies a nearline hard disk drive or an external MDisk for the newly discovered or external volume.

enterprise

Specifies an enterprise hard disk drive or an external MDisk for the newly discovered or external volume.

Note: This action applies to both copies if you are creating mirrored volume with two image mode copies by using this command.

-easytier on | off

Determines if the IBM Easy Tier function is allowed to move extents for this volume.

Note: The **-easytier** parameter must be followed by either *on* or *off*:

- If set to *on*, then Easy Tier functions are active.
- If set to *off*, then Easy Tier functions are inactive.

If the Easy Tier feature is enabled, and if a volume copy is striped and not being migrated, the following table applies:

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 739)
Off	One	On	inactive (see note 2 on page 739)
Off	Two	Off	inactive (see note 2 on page 739)
Off	Two	On	inactive (see note 2 on page 739)
Measure	One	Off	measured (see note 3 on page 739)
Measure	One	On	measured (see note 3 on page 739)
Measure	Two	Off	measured (see note 3 on page 739)

Storage pool Easy Tier setting	Number of tiers in the storage pool	Volume copy Easy Tier setting	Volume copy Easy Tier status
Measure	Two	On	measured (see note 3)
Auto	One	Off	measured (see note 3)
Auto	One	On	balanced (see note 4)
Auto	Two	Off	measured (see note 3)
Auto	Two	On	active (see note 5)
On	One	Off	measured (see note 3)
On	One	On	balanced (see note 4)
On	Two	Off	measured (see note 3)
On	Two	On	active (see note 5)

Note:

1. If the volume copy is in image or sequential mode or is being migrated then the volume copy Easy Tier status is 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 balanced, the Easy Tier function enables performance-based pool balancing for that volume copy.
5. When the volume copy status is active, the Easy Tier function operates in automatic data placement mode for that volume.
6. 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 except pool performance balancing are disabled for storage pools with a single tier, and that automatic data placement mode are enabled for all striped volume copies in a storage pool with two or more tiers.

Description

This command creates a new volume object. You can use the command to create various types of volume objects, making it one of the most complex commands.

Remember: You can create a striped volume only in a child pool - you cannot create sequential or image volumes in a child pool.

You must decide which storage pool or storage pools provide the storage for the volume. Use the **lsmdiskgrp** command to list the available storage pools and the amount of free storage in each pool. If you are creating a volume with more than one copy, each storage pool that you specify must have enough space for the size of the volume.

Important: The extent size for the storage pool can limit volume size. Consider the maximum volume size that you want to use when creating storage pools. Refer to the information on creating storage pools for a comparison of the maximum volume capacity for each extent size. The maximum is different for thin-provisioned volumes.

Choose an I/O group for the volume. This action determines which nodes in the system process the I/O requests from the host systems. If you have more than one I/O group, ensure that you distribute the volumes between the I/O groups so that the I/O workload is shared evenly between all nodes. Use the **lsiogr** command to show the I/O groups and the number of volumes that are assigned to each I/O group.

Note: It is normal for systems with more than one I/O group to have storage pools that have volumes in different I/O groups. FlashCopy processing can make copies of volumes whether the source and target volumes are in the same I/O group. If, however, you plan to use intra-system Metro or Global Mirror operations, ensure that both the master and auxiliary volume are in the same I/O group. The command returns the IDs of the newly created volume.

An encryption key cannot be used when creating an image mode MDisk. To use encryption (when the MDisk has an encryption key), the MDisk must be self-encrypting.

Specify the virtualization type by using the **-vtype** parameter; the supported types are sequential (*seq*), *striped*, and *image*.

sequential (seq)

This virtualization type creates the volume that uses sequential extents from the specified MDisk (or MDisk, 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 storage pool are used to create the volume. The striping is at an extent level; one extent from each managed disk in the group is used. For example, a storage pool with 10 managed disks uses one extent from each managed disk. It then 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 storage pool. 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 MDisk 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 volumes to be created when a managed disk already has data on it, perhaps from a previrtualized subsystem. When an image mode volume is created, it directly corresponds to the (previously unmanaged) managed disk that it was created from. Therefore, except for thin-provisioned image mode volumes, volume 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 system. After it is under the control of the system, you can migrate the volume from the single managed disk. When it is migrated, the volume is no longer an image mode volume.

You can add image mode volumes to an already populated storage pool with other types of volumes, such as a striped or sequential.

Important: An image mode volume must be 512 bytes or greater. At least one extent is allocated to an image mode volume.

Remember: If you create a mirrored volume from two image mode MDisk without specifying a **-size** value, the capacity of the resulting volume is the smaller of the two MDisk, and the remaining space on the larger MDisk is not accessible.

Attention:

1. Do not create a volume in an offline I/O group. You must ensure that the I/O group is online before you create a volume to avoid any data loss. This action applies in particular to re-creating volumes that are assigned the same object ID.
2. To create an image mode disk, you must already have a quorum disk present in the system 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 volumes per I/O Group or 8192 volume copies per system is reached.

The rate at which the volume copies resynchronize after loss of synchronization can be specified by using the **-syncrate** parameter. This table provides the relationship of the *syncrate* value to the data copied per second.

Note: These settings also affect the initial rate of formatting.

Table 117. Relationship between the syncrate value and the data copied per second

User-specified <i>syncrate</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

```
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 volume

```
mkvdisk -mdiskgrp Group0
-iogrp 0 -vtype image -mdisk mdisk2 -node 1
```

The resulting output:

```
Virtual Disk, id [2], successfully created
```

An invocation example

An invocation example for creating a new volume

```
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 thin-provisioned volume

```
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 compressed volume copy

```
mkvdisk -mdiskgrp 0 -iogrp 0 -size 1 -unit tb -rsize 0 -autoexpand -warning 0 -compressed
```

The resulting output:

```
Virtual Disk id [1], successfully created
```

An invocation example for creating a mirrored image-mode volume

```
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 volume

```
mkvdisk -iogrp 0 -mdiskgrp 0:1 -size 500 -copies 2
```

The resulting output:

```
Virtual Disk id [5], successfully created
```

An invocation example for configuring a mirror write algorithm priority

```
mkvdisk -mdiskgrp Group0 -iogrp 0 -vtype striped -mirrorwritepriority redundancy -size 500
```

The resulting output:

```
Virtual Disk id [5], successfully created
```

An invocation example to create a disk with default grain size

```
mkvdisk -iogrp 0 -mdiskgrp 0 -size 100 -rsize 5%
```

The resulting output:

```
Virtual Disk id [5], successfully created
```

An invocation example for creating a volume with I/O groups 0 and 1 in its I/O group access set

```
mkvdisk -iogrp 0 -mdiskgrp 0 -size 500 -accessiogrp 0:1
```

The resulting output:

```
Virtual Disk id [5], successfully created
```

An invocation example for creating a volume with warning considerations

```
mkvdisk -mdiskgrp 6 -size 200 -rsize 50 -iogrp 0
```

The resulting output:

```
Virtual Disk, id [2], successfully created
```

```
...
```

```
lsvdisk 2
```

```
...
```

```
warning 20 # threshold in MB = 50 x 80 / 100 = 40 MB; threshold as %age of volume capacity = 40 / 200 * 100 = 20
```

```
...
```


An invocation example for creating a volume with warning considerations

```
mkvdisk -mdiskgrp 6 -size 200 -rsize 50 -iogrp 0 -warning 80%
```

The resulting output:

```
Virtual Disk, id [2], successfully created
...
lsvdisk 2
...
warning 80    # displayed as %age of volume capacity
...
```

An invocation example for creating a volume with warning considerations

```
mkvdisk -mdiskgrp 6 -size 200 -rsize 50 -iogrp 0 -autoexpand
```

The resulting output:

```
Virtual Disk, id [2], successfully created
...
lsvdisk 2
...
warning 80    # displayed as %age of volume capacity
...
```

An invocation example to create a volume with the read cache enabled

```
mkvdisk -iogrp 0 -size 10 -unit gb -mdiskgrp 0 -cache readonly
```

The resulting output:

```
Virtual Disk, id [2], successfully created
```

An invocation example to create volume Group0

```
mkvdisk -mdiskgrp Group0 -iogrp io_grp0 -vtype image -mdisk 13 -node 1 -udid 1234 -tier tier_nearline
```

The resulting output:

```
Virtual Disk, id [0], successfully created
```

An invocation example to turn off formatting while creating volume Chelsea1

```
mkvdisk -mdiskgrp Chelsea1 -iogrp 0 -mdiskgrp 0:1 -size 500 -nofmtdisk -copies 2
```

The resulting output:

```
Virtual Disk, id [0], successfully created
```

mkvdiskhostmap

Use the **mkvdiskhostmap** command to create a new mapping between a volume and a host, which makes the volume accessible for input/output (I/O) operations to the specified host.

Syntax

```
➤ mkvdiskhostmap — [ -force ] — -host [ host_id ] [ host_name ] [ -scsi — scsi_num_arg ]
➤ [ vdisk_name ] [ vdisk_id ]
```

Parameters

-force

(Optional) Allows multiple volume-to-host assignments, which are not normally allowed.

-host *host_id* | *host_name*

(Required) Specifies the host to map the volume to, either by ID or by name.

-scsi *scsi_num_arg*

(Optional) Specifies the Small Computer System Interface (SCSI) logical unit number (LUN) ID to assign to this volume on the given host. The *scsi_num_arg* parameter contains the SCSI LUN ID that is assigned to the volume on the given host for all I/O groups that provide access to the volume. You must check your host system for the next available SCSI LUN ID on the given host bus adapter (HBA). If you do not specify the **-scsi** parameter, the next available SCSI LUN ID in each I/O group that provides access is provided to the host.

vdisk_name | *vdisk_id*

(Required) Specifies the name of the volume that you want to map to the host, either by ID or by name.

Description

This command creates a new mapping between the volume and the specified host. The volume 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 volume.

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 volumes 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 0, and so on. You can control the order in which the HBA discovers volumes 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.

The **mkvdiskhostmap** command fails if the:

- Host to which this mapping is being made is not associated with any one of the I/O groups in the volume access set
- Volume has more than one I/O group in its access set and the host being mapped to the volume does not support volumes being mapped from multiple I/O groups

Remember: iSCSI hosts can access volumes that are accessible through multiple I/O groups (as well as single I/O groups).

If you generate different SCSI LUN IDs, only one is returned. The returned ID is for the highest-numbered I/O group to which the volume was mapped. To view other values, issue **lshostvdiskmap** or **lsvdiskhostmap**.

The SCSI LUN ID is used for the highest numbered I/O group to which the volume is mapped.

Some HBA device drivers stop when they find a gap in the SCSI LUN IDs. For example:

- Volume 1 is mapped to Host 1 with SCSI LUN ID 1
- Volume 2 is mapped to Host 1 with SCSI LUN ID 2
- Volume 3 is mapped to Host 1 with SCSI LUN ID 4

When the device driver scans the HBA, it must stop after identifying volumes 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 volume assignments (assigning the same volume to multiple hosts, for example, which might be particularly useful for clustered system hosts assigning a volume to multiple hosts). Normally, multiple volume-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, a volume must be mapped to more than one host. This includes the IBM SAN File System. To map to more than one host, you must use the **mkvdiskhostmap** command with the **-force** parameter. For example:

```
mkvdiskhostmap -host host1 -force 4
mkvdiskhostmap -host host2 -force 4
```

Note: When assigning the same volume to multiple hosts, you should use the same SCSI ID for all hosts.

These commands create two host-to-volume mappings for volume 4 that map to host1 and host2. Omitting the **-force** parameter causes the mapping to fail if that volume 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 volume.

If a new mapping is created between a host (type `hide_secondary`) and a volume that is a secondary volume in a remote copy relationship, the:

- Mapping is created for configuration purposes (it can be changed or deleted)
- Secondary volume is not presented to the host

The mapped volume is presented to the host if the:

- Host type is changed to a type other than `hide_secondary`
- Remote copy relationship is stopped by specifying **-access**
- Volume is no longer a secondary volume because the remote copy relationship is deleted or switched

Note: You cannot specify this command if the volume is an auxiliary volume in an active-active relationship or if a volume is a change volume in any type of relationship.

An invocation example

```
mkvdiskhostmap -host host1 -scsi 1 5
```

The resulting output:

```
Virtual Disk to Host map, id [1], successfully created
```

mkvolume

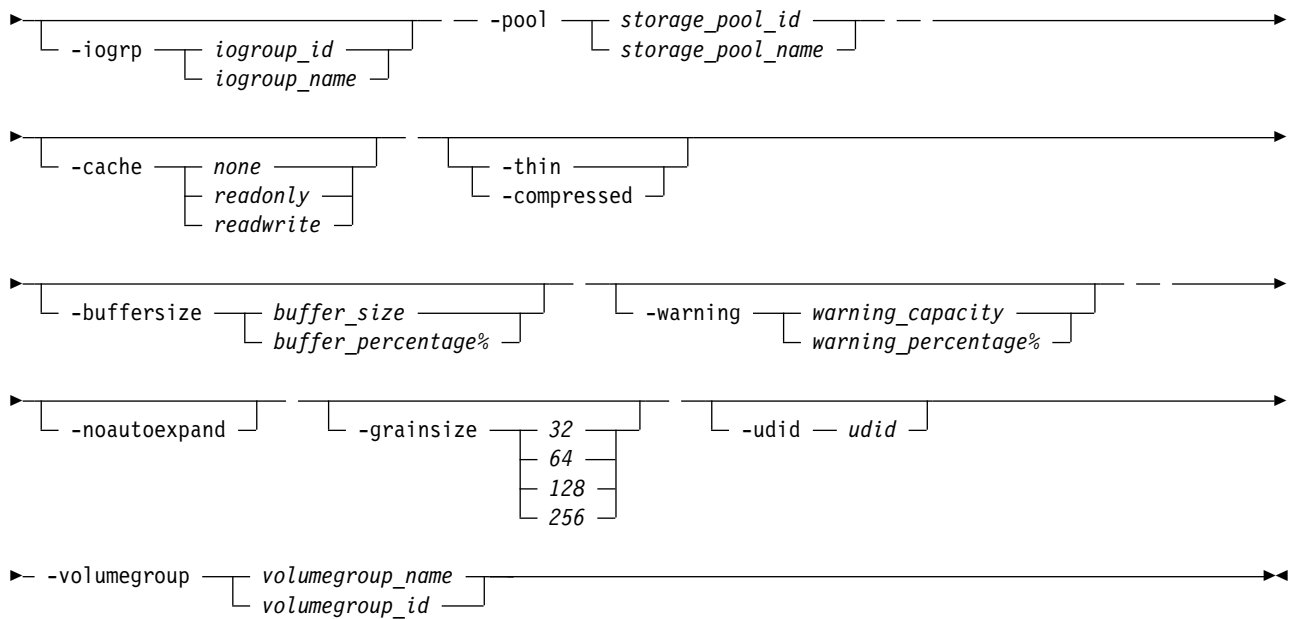
Use the **mkvolume** command to create an empty volume from existing storage pools. You can use this command for high availability configurations that include HyperSwap or stretched systems, but it can also be used for volumes that are not high availability.

Syntax

```

▶▶ mkvolume — [ name — name ] — size — disk_size — [ -unit — b | kb | mb | gb | tb | pb ] —

```



Parameters

-name *name*

(Optional) Specifies the name that is used for the volume that is created. This value must be an alphanumeric string 1 - 63 characters.

Remember: If you do not specify **-name**, a unique default name such as `volume1` is used.

-size *disk_size*

(Required) Specifies the capacity of the volume, which is used with the value of the unit. The default capacity is in MB. When the unit of bytes is used, all capacities must be in multiples of 512 bytes. An entire extent is reserved even if it is only partially used.

-unit *b | kb | mb | gb | tb | pb*

(Optional) Specifies the data units to use with the capacity that is specified by the **-size** parameter. The default unit type is *mb*.

-iogrp *iogroup_id_list | iogroup_name_list*

(Optional) Specifies the I/O group that the new volume is cached in. The value can be a colon-separated list of up to two I/O group IDs or names. If no value is specified, the caching I/O group is selected based on the storage pool site. If you do not specify the **-iogrp** parameter, the caching I/O group is selected by the system.

Important: If two I/O groups are specified, they must be in different sites, and the storage pools that are specified must be in different sites. The order of the sites must correspond.

If you create a HyperSwap volume, the caching I/O groups are selected based on the sites of the storage pools.

-pool *storage_pool_id_list | storage_pool_name_list*

(Required) Specifies the storage pool in which to create the new volume. The value must be a colon-separated list of up to two storage pool IDs or names.

Note: If one storage pool is specified, a basic volume is created with one copy.

On systems with standard topology, a mirrored volume can be created by specifying two storage pools.

On systems with a stretched or hyperswap topology, a highly available volume can be created by specifying two storage pools in different sites.

-cache *none* | *readonly* | *readwrite*

(Optional) Specifies the caching options for the volume. Use one of the following valid entries:

- *readwrite* enables the cache for the volume (default)
- *readonly* disables write caching but allows read caching for a volume
- *none* disables the cache mode for the volume

-thin

(Optional) Specifies that the volume is to be created with thin-provisioning. You cannot specify this parameter with **-compressed**. If you do not specify **-thin** or **-compressed**, the volume that is created is fully allocated.

-compressed

(Optional) Specifies that the volume is to be created compressed. If the **-iogrp** parameter is not specified, the least used I/O group is used for compressed copies (considering the subset of I/O groups that support compression).

Remember: This command fails if no I/O groups support compression. If there are two sites, both sites must have at least one I/O group that supports compression.

You cannot specify this parameter with **-thin**. If you do not specify **-thin** or **-compressed**, the volume that is created is fully allocated.

-buffersize *buffer_size* | *buffer_percentage*

(Optional) Specifies the pool capacity the volume attempts to reserve as a buffer for thin-provisioned and compressed volumes. You must specify either **-thin** or **-compressed** with this parameter. The default value is 2%.

-warning *warning_capacity* | *warning_percentage*

(Optional) Specifies a threshold at which a warning error log is generated for volumes. A warning is generated when the used disk capacity on the thin-provisioned volume exceeds the specified threshold. You must specify either **-thin** or **-compressed** with this parameter. The default value is 80%.

-noautoexpand

(Optional) Specifies that the volume not automatically expand as it is written to. The available buffer capacity decreases as the used capacity increases. The volume copy goes offline if the buffer capacity is fully used. The buffer capacity can be increased by specifying **expandvdiskspace -rsize**. You must specify either **-thin** or **-compressed** with this parameter. If you do not specify **-noautoexpand**, the volume automatically expands as it is written to.

-grainsize *32* | *64* | *128* | *256*

(Optional) Sets the grain size (KB) for a thin-provisioned volume. If you are using the thin-provisioned volume in a FlashCopy map, use the same grain size as the map grain size for best performance. If you are using the thin-provisioned volume directly with a host system, use a small grain size. The grain size value must be 32, 64, 128, or 256 KB. The default is 256 KB.

-udid *udid*

(Optional) Specifies the unit number *udid* for the volume.

Important: The *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 through 32767, or a hexadecimal number from 0 through 0x7FFF. A hexadecimal number must be preceded by 0x (for example, 0x1234).

-volumegroup *volume_group_name* | *volume_group_id*

(Optional) Specifies the volume group that a volume belongs to. The value must be an alphanumeric string for the volume group name and the value must be a number for the volume group ID.

Description

This command creates an empty volume, which is a formatted (zeroed) volume, by using storage from existing storage pools. You can also create a highly available volume on systems with stretched or hyperswap topology.

Use the **mkimagevolume** command to create a new volume by importing existing data on a managed disk.

An invocation example to create a volume in storage pool 0

```
mkvolume -pool 0 -size 1000
```

The detailed resulting output:

Volume, id [0], successfully created.

An invocation example for creating a thin-provisioned stretched volume on a system with stretched topology

```
mkvolume -pool site1pool:site2pool -size 1 -unit tb -thin
```

The detailed resulting output:

Volume, id [1], successfully created.

An invocation example to create a HyperSwap volume with a hyperswap topology

```
mkvolume -pool site1pool:site2pool -size 200
```

The detailed resulting output:

Volume, id [2], successfully created.

An invocation example to assign a volume group

```
mkvolume -volume group 1
```

The detailed resulting output:

Volume, id [3], successfully created.

mkvolumegroup

Use the **mkvolumegroup** command to create and configure a new volume group.

Syntax

```
➤— mkvolumegroup — [ -name — volume group_name ] —➤
```

Parameters

-name *volume group_name*

(Optional) Specifies a volume group name. The value must be an alphanumeric value. If you do not specify a volume group name, one is automatically created and assigned to the volume group.

Description

This command creates and configures a new volume group.

An invocation example

```
mkvolumegroup
```

The resulting output:

```
Volume Group, id [0], successfully created
```

An invocation example

```
mkvolumegroup -name Sunday
```

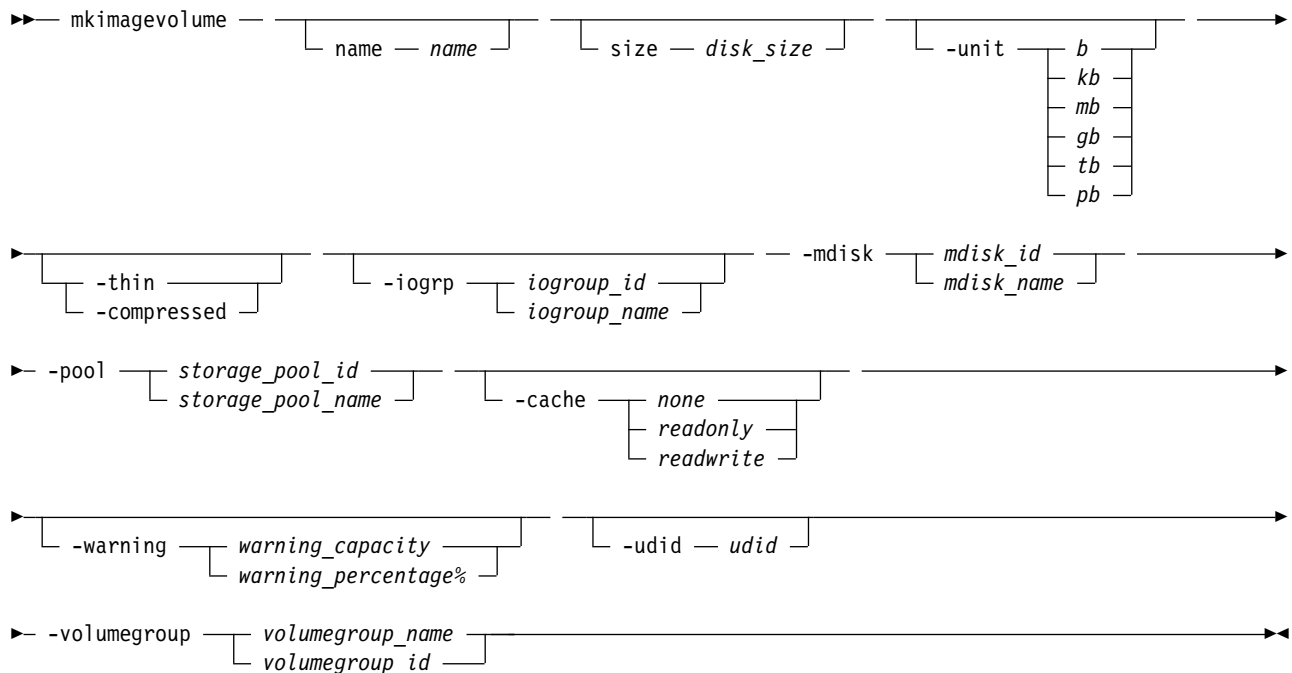
The resulting output:

```
Volume Group, id [1], successfully created
```

mkimagevolume

Use the **mkimagevolume** command to create an image mode volume by importing (preserving) data on a managed disk from another storage system.

Syntax



Parameters

-name *name*

(Optional) Specifies the name that is used for the volume that is created. This value must be an alphanumeric string 1 - 63 characters.

Remember: If you do not specify **-name**, a unique default name such as `volume1` is used.

-size *disk_size*

(Required if **-thin** or **-compressed** is specified) Specifies the capacity of the volume, which is used with the value of the unit. The default capacity is in MB. When the unit of bytes is used, all capacities must be in multiples of 512 bytes. An entire extent is reserved even if it is only partially used.

Remember: This parameter is optional if **-thin** or **-compressed** are not specified. If this parameter is not specified the volume is fully allocated.

For thin and compressed volumes, the real capacity is set from the MDisk size.

-unit *b | kb | mb | gb | tb | pb*

(Optional) Specifies the data units to use with the capacity that is specified by the **-size** parameter. The default unit type is *mb*.

-thin

(Optional) Specifies that the volume is to be created with thin-provisioning. You cannot specify this parameter with **-compressed**. If you do not specify **-thin** or **-compressed**, the volume that is created is fully allocated.

-compressed

(Optional) Specifies that the volume is to be created compressed. If the **-iogrp** parameter is not specified, the least used I/O group is used for compressed copies (considering the subset of I/O groups that support compression).

Remember: This command fails if no I/O groups support compression.

You cannot specify this parameter with **-thin**. If you do not specify **-thin** or **-compressed**, the volume that is created is fully allocated.

-iogrp *iogroup_id | iogroup_name*

(Optional) Specifies the I/O group that the new volume is cached in.

-mdisk *mdisk_id mdisk_name*

(Required) Specifies which currently unused MDisk to use to create the image mode volume.

-pool *storage_pool_id | storage_pool_name*

(Required) Specifies the storage pool in which to create the new volume. The value for *storage_pool_id* must be a numerical value.

-cache *none | readonly | readwrite*

(Optional) Specifies the caching options for the volume. Valid entries are:

- *readwrite* enables the cache for the volume (default)
- *readonly* disables write caching but allows read caching for a volume
- *none* disables the cache mode for the volume

-warning *warning_capacity | warning_percentage*

(Optional) Specifies a threshold at which a warning error log is generated for volume copies. A warning is generated when the used disk capacity on the thin-provisioned volume exceeds the specified threshold. You must specify either **-thin** or **-compressed** with this parameter. The default value is 80%.

-udid *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 from 0 through 32767, or a hexadecimal number from 0 through 0x7FFF. A hexadecimal number must be preceded by 0x (for example, 0x1234).

Remember: When you create a HyperSwap volume, this value is set only on the master volume.

-volumegroup *volumegroup_name | volumegroup_id*

(Optional) Specifies the volume group that a volume image belongs to. The value must be an alphanumeric string for the volume group name and the value must be a number for the volume group ID.

Description

Use the **mkimagevolume** command to create a new image mode volume. This command is used to import a volume, preserving existing data.

Import a fully allocated image mode volume into storage pool 0 with MDisk 2 at full capacity

```
mkimagevolume -mdisk 2 -pool 0
```

The detailed resulting output:

Volume, id [0], successfully created.

Import a thin-provisioned image mode volume (with virtual capacity 25GB) into storage pool 1 with MDisk 7

```
mkimagevolume -mdisk 7 -pool 1 -thin -size 25 -unit gb
```

The detailed resulting output:

Volume, id [2], successfully created.

Assign a volume group to a volume image

```
mkimagevolume -volume group 1
```


The detailed resulting output:

Volume, id [3], successfully created.

movevdisk

Use the **movevdisk** command to move the preferred node of a volume either within the same caching I/O group or to another caching I/O group.

Syntax



Parameters

-iogrp *iogrp_id* | *iogrp_name*

(Optional) Specifies the I/O group to move the volume to.

-force

(Optional) Use the force parameter to force the volume to be removed from an I/O group. This option overrides the cache flush mechanism.

Remember:

- If you specify the **-force** parameter the contents of the cache are discarded and the volume might be corrupted by the loss of the cached data. Use the **-force** parameter with caution.
- If the force parameter is used to move a volume that has out-of-sync copies, a full resynchronization is required.

-node *node_id* | *node_name*

(Optional) Specifies the node ID or name assigned as the preferred node.

vdisk_id | *vdisk_name*

(Required) Specifies the volume to move.

Description

Use the **movevdisk** to migrate a single volume to a new I/O group - repeat this action for other volumes as required. This command can also move the preferred node of a volume without changing the caching I/O group, but it does not change which I/O groups can access the volume (only the caching I/O group is changed).

Important: You cannot migrate or move an image mode volume.

A compressed volume can also be moved, and you can specify the preferred node in the new I/O group. You can move a volume that is in a Flash Copy (FC) mapping, but the FC bitmaps remain in the original I/O group. You cannot move the volumes when the FC mapping is in preparing or prepared state. Additionally, a volume can be moved if it is the target of a FC mapping that is in stopping state.

You cannot move a volume to change the caching I/O group for a volume that is in a Global Mirror, Metro Mirror, or HyperSwap relationship, regardless of whether it is a primary, secondary, or change volume. To move a volume in a Global Mirror, Metro Mirror, or HyperSwap relationship, the relationship must first be deleted. You can change a preferred node without changing caching I/O group for this type of a volume.

If the volume is offline, use one of the **recovervdisk** commands to recover the volume and bring it back online. To specify a preferred node for the volume, use the **-node node_id | node_name** parameter with the **movevdisk** command. Use the **movevdisk** command to change the I/O group with which this volume is associated.

Important: Do not move:

- A volume to an offline I/O group under any circumstance;

Remember: To avoid data loss, make sure that the I/O group is online before moving the volume.

- An offline volume to the recovery I/O group

You can migrate a volume to a new I/O group to manually balance the workload across the nodes in the clustered system. You might end up with a pair of nodes that are overworked and another pair that are underworked.

Remember: You cannot move a volume if that volume is being formatted.

If the volume is a target of a FlashCopy mapping with a source volume in an active-active relationship then the new I/O group must be in the same site as the source volume. The system only permits moving a volume in a remote copy relationship if the move does not change the I/O group (it changes the preferred node). If the volume is in an active-active relationship the new I/O group must be in the same site as the source I/O group.

Note: Remote copy includes Metro Mirror, Global Mirror, and HyperSwap.

An invocation example to move DB_Volume to I/O group 2

```
movevdisk -iogrp 2 DB_Volume
```

The resulting output:

No feedback

An invocation example to move DB_Volume to I/O group IOGRP3 with a new preferred node id 7

```
movevdisk -iogrp IOGRP3 -node 7 DB_Volume
```

The resulting output:

No feedback

An invocation example to change preferred node of volume DB_Volume with a new preferred node id as 8 in the same IOGRP

```
movevdisk -node 8 DB_Volume
```

The resulting output:

No feedback

recovervdisk

Use the **recovervdisk** command to acknowledge volume data loss and brings the volume back online.

Syntax

```
➤—recovervdisk— [—copy— —copy_id—] [vdisk_name | vdisk_id]—➤
```

Parameters

vdisk_name | *vdisk_id*

(Required) Specifies the volume to recover.

-copy *copy_id*

(Optional) Specifies the ID of the copy to recover.

Description

The specified volume, and all copies if mirrored, are recovered and brought back online. If the volume is thin-provisioned or has thin-provisioned copies, this command triggers the thin-provisioned repair process. If the volume is mirrored, the **recovervdisk** command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored using the **lsdisksyncprogress** command. The volume remains online during the resynchronization process.

The **recovervdisk** command also starts the repair of any thin-provisioned copies that have a `fast_write_state` of `corrupt`. The progress of the repair process can be monitored using the **lsrepairsevdiskcopyprogress** command.

A volume that is still offline because it is being repaired following the **recovervdisk** command has a `fast_write_state` of `repairing`. The volume is brought online when the repair process is complete.

An invocation example (to recover volume 45)

```
recovervdisk vdisk45
```

An invocation example (to recover copy 0 of volume 45) :

```
recovervdisk -copy 0 vdisk45
```

recovervdiskbycluster (Discontinued)

Attention: The **recovervdiskbycluster** command has been discontinued. Use the **recovervdiskbysystem** command instead.

recovervdiskbyiogrp

Use the **recovervdiskbyiogrp** command to acknowledge data loss for all volumes in the specified I/O group with a `fast_write_state` of corrupt and brings the volumes back online.

Syntax

```
►►—recovervdiskbyiogrp—  
```

Parameters

io_group_name | *io_group_id*
(Required) Specifies the I/O group for volume recovery.

Description

All volumes 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 volumes are thin-provisioned or have thin-provisioned copies, the **recovervdiskbyiogrp** command triggers the thin-provisioned repair process. If volumes are mirrored, the command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored by using the **lsrepairsevdiskcopyprogress** command. volumes remain online during the resynchronization process.

If none of the volumes 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 volumes. The progress of the repair process can be monitored using the **lsrepairsevdiskcopyprogress** command. If there are no corrupt volumes or no repairs to copies are required, no error is returned.

Volumes that are still offline because they are being repaired following the **recovervdiskbyiogrp** command have a `fast_write_state` of repairing. Volumes are brought online when the repair process is complete.

An invocation example

```
recovervdiskbyiogrp iogrp2
```


The resulting output:

No feedback

recovervdiskbysystem

Use the **recovervdiskbysystem** command to acknowledge data loss for all volumes in the clustered system (system) with a `fast_write_state` of corrupt and bring the volumes back online.

Syntax

```
►►—recovervdiskbysystem— 
```

Parameters

There are no parameters for this command.

Description

All volumes in the system that have a `fast_write_state` of corrupt; and all copies, if mirrored, are recovered and brought back online. If any of the volumes are thin-provisioned or have thin-provisioned copies, the **recovervdiskbysystem** command triggers the thin-provisioned repair process. If volumes are mirrored, the command triggers a resynchronization from a synchronized copy. The progress of the resynchronization can be monitored by using the **lsvdisksyncprogress** command. Volumes remain online during the resynchronization process.

If none of the volumes in the system have a `fast_write_state` of corrupt, the **recovervdiskbysystem** command still starts the repair process for any corrupt copies of mirrored volumes. The progress of the repair process can be monitored using the **lsrepairsevdiskcopyprogress** command. If there are no corrupt volumes or no repairs to copies are required, no error is returned.

Volumes that are still offline because they are being repaired following the **recovervdiskbysystem** command have a `fast_write_state` of repairing. Volumes are brought online when the repair process is complete.

An invocation example

```
recovervdiskbysystem
```

The resulting output:

No feedback

repairsevdiskcopy

The **repairsevdiskcopy** command repairs the metadata on a thin-provisioned volume copy or a compressed volume copy.

Syntax

```
➤➤—repairsevdiskcopy— [ -copy — 0 | 1 ] [ vdisk_name | vdisk_id ] ➤➤
```

Parameters

-copy 0 | 1

(Optional) Specifies the volume copy to repair.

vdisk_name | vdisk_id

(Required) Specifies the volume to repair.

Description

The **repairsevdiskcopy** command repairs the metadata on a thin-provisioned volume or compressed volume copy. Run this command only when you are directed by the fix procedures or by your product support information.

Running the command automatically detects corrupted metadata. The command holds the 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 volume 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

```
repairsevdiskcopy vdisk8
```

The resulting output:

No feedback

repairvdiskcopy

Use the **repairvdiskcopy** command to detect and (optionally) correct any volume copies that are not identical.

Syntax

```
➤—repairvdiskcopy— [ -medium | -resync | -validate ] [ -startlba lba ] [ vdisk_name | vdisk_id ] ➤
```

Parameters

-medium

(Optional) Converts sectors that contain different readable data into virtual medium errors on the specified volume. It fixes preexisting medium errors found on only one volume copy by replacing them with data from the other volume copy. This parameter cannot be used with the **-validate** and **-resync** parameters. You must specify one of the three parameters.

-resync

(Optional) Corrects sectors that contain different readable data by copying contents from the primary volume copy to other copies on the specified volume. It fixes preexisting medium errors found on only one volume by replacing them with data from the other volume. This parameter cannot be used with the **-medium** and **-validate** parameters. You must specify one of the three parameters.

-validate

(Optional) Reports the first difference in readable data found on synchronized online copies of the specified volume, 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. The LBA must be specified in hex, with a 0x prefix.

vdisk_name | *vdisk_id*

(Required) Specifies the volume to repair. You must specify this parameter last on the command line.

Description

The **repairvdiskcopy** command detects and optionally, corrects any volume copies that are not identical. For the purposes of comparison, preexisting medium errors found on only one volume are ignored and fixed by replacing them with data from the other volume copy. The results are logged to the SAN Volume Controller error log. The **-validate** parameter compares synchronized online copies of the specified volume. 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 volume copies. You cannot use this command with a volume that is fast formatting.

You must specify **-validate**, **-medium**, or **-resync**.

Attention:

1. Before you run the **repairvdiskcopy** command, ensure that all volume copies are synchronized.
2. Only one **repairvdiskcopy** command can run on a volume 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 volume 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 to 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 volume copies are not identical.

During **repairvdiskcopy** command operation, the volume 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 volume that is being repaired. To suspend the repair process, set the synchronization rate of the volume to 0 using the **chvdisk** command.

An invocation example

```
repairvdiskcopy -resync -startlba 0x0 vdisk8
```


The resulting output:

No feedback

restorevolume

Use the **restorevolume** command to restore a volume from a snapshot generation.

Syntax



Parameters

-fromuid *volume_UID*

(Optional) Specifies the volume snapshot to restore (specified by volume UID). The value must be a number.

Use this parameter to restore a snapshot from a different volume. This means that the specified UID must be different from the UID of the volume being restored (when you specify *volume_name* or *volume_id*).

Note: The volume being restored to cannot have cloud snapshot enabled if you specify this parameter.

-restoreuid

(Optional) Specifies the UID of the restored volume should be set to the UID of the volume snapshot that is being restored. You must specify **-fromuid** with this parameter.

-generation *gen_id*

(Optional) Specifies the snapshot generation to restore. The value must be a number.

-deletelatergenerations

(Optional) Specifies that all snapshot generations should be deleted after the generation is restored.

Note: This parameter is required when the volume has cloud snapshot enabled and the generation being restored is not the most recent snapshot of the volume.

-cancel

(Optional) Cancels the restore operation.

volume_name | *volume_id*

(Required) Specifies the volume name or ID to restore. The value for the volume ID must be a number and the value for the volume name must be an alphanumeric string.

Description

This command restores a volume from a snapshot generation.

The restore operation is performed directly on the volume specified by volume name or volume ID (without use of a temporary volume). The volume is offline while the restore operation is in progress. If the restore process is canceled before it completes the data on the volume is inconsistent and not usable.

An invocation example

To restore an earlier generation (generation 3) for volume `volume7`:

```
restorevolume -generation 3 -deletelatergenerations volume7
```

The resulting output:

No feedback

An invocation example

To restore the most recent snapshot (generation 5) for volume ID 7:

```
restorevolume -generation 5 volume7
```

The resulting output:

No feedback

An invocation example

```
restorevolume -generation 1 -fromuid 6005076400B70038E00000000000001C 1
```


The resulting output:

No feedback

rmvdisk

Use the **rmvdisk** command to delete a volume. This command cannot be used for high availability volumes. Use the **rmvolume** command for high availability volumes.

Syntax



Parameters

-force

(Optional) Deletes the specified volume, even if mappings still exist between this volume and one or more hosts. This parameter deletes any host-to-volume mappings and any FlashCopy mappings that exist for this volume.

Important: To prevent an active volume from being deleted unintentionally, you can use a global system setting to enable volume protection (see the **chsystem** command). You can specify a time period for which the volume must be idle before you can delete it. If volume protection is enabled and the time period has not expired, the volume deletion fails even if the **-force** parameter is used." If the **-force** deletion of a volume causes dependent mappings to be stopped, any target volumes for those mappings that are in Metro Mirror or Global Mirror relationships are also stopped. The dependent mappings can be identified by using the **lsvdiskdependentmaps** command on the volume that you want to delete.

Note: Using the **-force** parameter might result in a data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of volume data loss.

If you do not specify this parameter a volume cannot be deleted while a backup operation is in progress. Additionally, a volume containing image mode copies cannot be deleted while a restore operation is in progress (if the volume contains inconsistent data).

-removehostmappings

(Optional) Removes all host mappings for the specified volume before deleting the volume itself.

Note: Using the **-removehostmappings** parameter might result in a data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of volume data loss.

vdisk_id | vdisk_name

Specifies the name of the volume to delete, either by ID or by name.

Note: To deactivate compression, use the **rmvdiskcopy** to delete the last compressed volume copy for an I/O group.

Description

This command deletes an existing managed mode volume or an existing image mode volume. The extents that made up this volume are returned to the pool of free extents that are available on the storage pool, if the volume is in managed mode.

Remember: If you issue this command, any data that was on the volume is lost. Before you issue this command, ensure that the volume (and any data that resides on it) is no longer required.

This command is unsuccessful if:

- Volume protection is enabled (using the **chsystem** command)
- The volume being removed has received I/O within the defined volume protection time period

With an active-active relationship, either or both of the master and auxiliary volumes can provide the information for host systems to read through the master volume ID. To remove the auxiliary volume from the relationship, delete the relationship so hosts can access the master copy.

Remember: Any FlashCopy mappings with the specified volume as their source volume are deleted when you specify this command.

Deleting a managed mode volume

When you use this command to delete a managed mode volume, all the data on the volume is deleted. The extents that make up the volume are returned to the pool of free extents that are available in the storage pool.

If host mappings exist for the volume, 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 volume as source or target are deleted, other mappings in a cascade might be stopped, and then the volume is deleted. The **-force** parameter also deletes any Metro Mirror or Global Mirror relationships that exist for the specified volume (and any information that is not staged in the fast write cache).

If the volume is in the process of migrating to an image mode volume (using the **migratetoimage** command), the deletion fails unless you use the **-force** parameter. If you use the **-force** parameter, the migration is halted and then the volume is deleted. Before you issue this command, ensure that the volume (and any data that resides on it) is no longer required.

Deleting an image mode volume

If the volume 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 volume before the volume is deleted. This process can take several minutes to complete, and is indicated by the *fast_write_state* state of the volume being empty. If the **-force** parameter is specified, the fast-write data is discarded and the volume 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, the system attempts to move the data out of the cache; this process can time out, however.

If there are any virtual medium errors on the volume, 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 volume, the deletion fails unless you use the **-force** parameter. If you use the **-force** parameter, mappings are deleted and the volume is deleted. If there is any data that is not staged in the fast write cache for this volume, the deletion of the volume fails. When the **-force** parameter is specified, any data that is not staged in the fast write cache is deleted. Deleting an image mode volume causes the managed disk that is associated with the volume to be removed from the storage pool. The mode of the managed disk is returned to unmanaged.

If the relationship is in *consistent_copying* or *consistent_stopped* state, and the change volume is being used by a Global Mirror relationship using multicycling mode, the relationship moves to *inconsistent_copying* or *inconsistent_stopped* state.

Note: If the relationship is part of a consistency group entire group is affected by this state transition. The secondary volume becomes corrupt, and inaccessible for host input/output I/O data if:

- A changed volume is part of an idling relationship
- The changed volume is being used for secondary protection
- The background copy process is still migrating the change volume data to the secondary volume

You must issue **recovervdisk** to gain access to the volume contents once more. If a change volume was part of an idling relationship and being used for Global Mirror relationship using multicycling mode, and the relationship was deleted but the background copy process continued and is still migrating data to the secondary volume then the secondary volume also becomes corrupt. In any of these cases, this **recovervdisk** fails without **-force** being specified.

Note:

- The **-force** parameter must be used if **rmvdisk** is specified and rejected if the volume is a change volume for a relationship.
- If the volume is a change volume for a relationship, specifying **rmvdisk** with **-force** removes the change volume from the relationship.

An invocation example

```
rmvdisk -force vdisk5
```

The resulting output:

No feedback

An invocation example

```
rmvdisk -removehostmappings vdisk3
```

The resulting output:

No feedback

rmmetadatavdisk

Use the **rmmetadatavdisk** command to detach a file system or remove a block device (that is based on the volume with owner type `host_integration_metadata`) in a configuration node.

Syntax

```
►►— rmmetadatavdisk — [ -ignorevvolsexist ] —►►
```

Parameters

-ignorevvolsexist

(Optional) Specifies that you want the system to delete the metadata volume (including volumes with owner type `vvol`).

Description

This command removes the metadata volume from a storage pool.

When **-ignorevvolsexist** is specified, only the metadata volume is deleted. Additionally, volumes with owner type `vvol` are not affected.

An invocation example

```
rmmetadavdisk -ignorevolsexist
```

The resulting output:

No feedback

rmvdiskcopy

Use the **rmvdiskcopy** command to remove a volume copy from a volume. This command cannot be used for high availability volumes.

Syntax

```
➤—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 volume, which deletes the entire volume. The parameter also forces the deletion of a nonmirrored volume, a copy that is migrating to image mode, or an image-mode copy that has virtual medium errors.

Important: To prevent an active volume from being deleted unintentionally, you can use a global system setting to enable volume protection (see the **chsystem** command). You can specify a time period for which the volume must be idle before you can delete it. If volume protection is enabled and the time period has not expired, the volume deletion fails even if the **-force** parameter is used."

vdisk_name | *vdisk_id*

(Required) Specifies the volume 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 volume. The command fails if all other copies of the volume are not synchronized; in this case, you must specify the **-force** parameter, delete the volume or more, or wait until the copies are synchronized.

Remember: This command is unsuccessful if:

- Volume protection is enabled
- The last volume copy being deleted has received I/O within the defined volume protection time period

These changes apply to this command only when deleting the last synchronized copy of a volume or removing the entire volume.

An invocation example

```
rmvdiskcopy -copy 1 vdisk8
```

The resulting output:

No feedback

rmvdiskaccess

Use the **rmvdiskaccess** command to delete one or more I/O groups from the set of I/O groups in which a volume can be made accessible to hosts.

Syntax

```
►►— rmvdiskaccess — — -iogrp ————  $\left[ \begin{array}{l} iogrp\_id\_list \\ iogrp\_name\_list \end{array} \right]$  ————  $\left[ \begin{array}{l} vdisk\_id \\ vdisk\_name \end{array} \right]$  —————►►
```

Parameters

-iogrp *iogrp_id_list* | *iogrp_name_list*

(Required) Specifies a list of I/O groups to remove from the I/O group access set of the volume.

vdisk_id | *vdisk_name*

(Required) Specifies the volume from which to remove access I/O groups.

Description

The **rmvdiskaccess** command removes I/O groups from the volume access set. However, it cannot remove all I/O groups from the access set; a volume must have at least one I/O group in an access set. When an I/O group is removed from the access set, all host mappings created through that I/O group (for the volume) are deleted. Consequently, you cannot access the volume through any related I/O group nodes.

Remember: If an I/O group in the list is not in the access set, no error is generated, but no action is taken for that I/O group.

An invocation example to remove I/O groups 2 and 3 from the volume access set for volume ID 3

```
rmvdiskaccess -iogrp 2:3 3
```

The resulting output:

No feedback

rmvdiskhostmap

Use the **rmvdiskhostmap** command to delete an existing host mapping the volume is no longer accessible for input/output (I/O) transactions on the given host.

Syntax

```
►►— rmvdiskhostmap — — -host ————  $\left[ \begin{array}{l} host\_id \\ host\_name \end{array} \right]$  ————  $\left[ \begin{array}{l} vdisk\_id \\ vdisk\_name \end{array} \right]$  —————►►
```

Parameters

-host *host_id* | *host_name*

(Required) Specifies the host that you want to remove from the map with the volume, either by ID or by name.

vdisk_id | *vdisk_name*

(Required) Specifies the name of the volume 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 volume and the host. This effectively stops the volume from being available for I/O transactions on the given host.

This command also deletes a Small Computer System Interface (SCSI) or persistent reservation that a host has on a volume. Once the reservation is removed, a new host is allowed to access the volume in the future because the original host no longer has access.

Note: The **rmvdiskhostmap** command deletes the host mapping for all I/O groups in the access I/O group set of the volume.

Use caution when you process this command because to the host, it seems as if the volume has been deleted or is offline.

Remember: This command is unsuccessful if:

- Volume protection is enabled
- The host mapping being deleted is mapped to any volume that has received I/O within the defined volume protection time period

An invocation example

```
rmvdiskhostmap -host host1 vdisk8
```

The resulting output:

No feedback

rmvolumegroup

Use the **rmvolumegroup** command to remove a volume group from a system.

Syntax

```
➤— rmvolumegroup — [ volume_group_name | volume_group_id ] —➤
```

Parameters

volume_group_name | *volume_group_id*

(Required) Specifies the volume group name or ID for the volume to remove from the system. The value for the volume group ID must be a number and the value for the volume group name must be an alphanumeric string.

Description

This command removes a volume group from a system.

Note: You cannot delete a volume group cannot if it contains active volumes.

An invocation example

```
rmvolumegroup Vardy1
```


The resulting output:

No feedback

rmvolumebackupgeneration

Use the **rmvolumebackupgeneration** command to delete a volume backup or cancel a volume snapshot operation that is in progress.

Syntax



Parameters

-volume *volume_name* | *volume_id*

(Optional) Specifies the volume snapshot by name or ID.

Note: The volume must exist on the local system.

The value for the volume name must be an alphanumeric string and the value for the volume ID must be a number. This parameter is mutually exclusive with **-uid**.

-uid *volume_UID*

(Optional) Specifies the volume snapshot UID. The value for a volume UID must be a number 0 - 32. This parameter is mutually exclusive with **-volume**.

-generation *gen_id*

(Optional) Specifies the snapshot generation to be deleted for the volume. Only a single snapshot generation is deleted. If the specified generation is for a snapshot operation that is in progress, that snapshot operation is canceled. The value for the generation ID must be a number. This command is mutually exclusive with **-all**.

-all

(Optional) Specifies deletion for all volume snapshots (which cancels all snapshot generations as well). This command is mutually exclusive with **-generation**.

Description

This command deletes a volume snapshot or cancels a volume snapshot operation that is in progress.

Note: If the:

- Command completes immediately, the delete operation is performed asynchronously
- Volume has more than one snapshot generation, you cannot delete the most recent snapshot generation

An invocation example

To delete snapshot generation 22 for volume with the UID 600507680CA880DF1800000000000007:

```
rmvolumebackupgeneration -uid 600507680CA880DF18000000000000007 -generation 22
```

The resulting output:

No feedback

An invocation example

To cancel the current snapshot generation 5 that is in progress for volume vdisk7:

```
rmvolumebackupgeneration -volume vdisk7 -generation 5
```

The resulting output:

No feedback

An invocation example

To remove all volume snapshots with the name `vdisk10`:

```
rmvolumebackupgeneration -volume vdisk10 -all
```

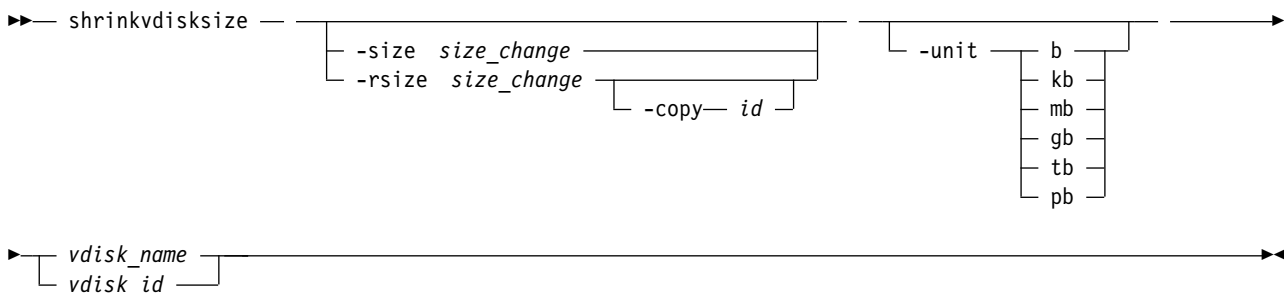
The resulting output:

No feedback

shrinkvdiskspace

Use the **shrinkvdiskspace** command to reduce the size of a volume by the specified capacity.

Syntax



Parameters

-size *size_change*

(Optional) Specifies the size reduction (change in size) for the designated volume. The **-size** parameter cannot be used with the **-rsize** parameter. You must specify either **-size** or **-rsize**.

Important: This parameter does reduce the size of a volume (the specified virtual size capacity).

-rsize *size_change*

(Optional) Reduces the real size of a thin-provisioned volume by the specified amount. This indicates the change in size as a result of the reduction. Specify the *size_change* value using an integer. Specify the units for a *size_change* integer using the **-unit** parameter; the default is MB. You must specify either **-rsize** or **-size**.

-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 volume are reduced. This parameter is required if the volume is mirrored and only one copy is thin-provisioned.

-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 volume that you want to modify, either by ID or by name.

Description

The **shrinkvdiskspace** command reduces the capacity that is allocated to the particular volume by the amount that you specify. You cannot shrink the real size of a thin-provisioned 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. You cannot use **shrinkvdisksize** if the volume is fast formatting.

The command can be used to shrink the physical capacity that is allocated to a particular volume by the specified amount. The command can also be used to shrink the virtual capacity of a thin-provisioned volume without altering the physical capacity assigned to the volume. To change the capacity of a non-thin-provisioned disk, use the **-size** parameter. To change the real capacity of a thin-provisioned disk, use the **-rsize** parameter. To change the virtual capacity of a thin-provisioned disk, use the **-size** parameter.

Volumes can be reduced in size, if required.

When the virtual size of a thin-provisioned volume 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 volume, all copies of the volume must be synchronized.

Attention: If the volume contains data that is being used, do not shrink the volume without backing up the data first.

The clustered system (system) arbitrarily reduces the capacity of the volume by removing a partial, one or more extents from those allocated to the volume. You cannot control which extents are removed and so you cannot assume that it is unused space that is removed.

Remember: Before you shrink a volume, validate that the volume is not mapped to any host objects.

You can determine the exact capacity of the source or master volume by issuing the **lsvdisk -bytes vdiskname** command. Shrink the volume by the required amount by issuing the **shrinkvdisksize -size size_change-unit b | kb | mb | gb | tb | pb vdisk_name | vdisk_id** command.

Remember:

1. You cannot resize (shrink) an image mode volume.
2. You cannot resize (shrink) the disk if the volume contains data.
3. You cannot resize (shrink) a volume that is part of a file system.
4. You cannot resize (shrink) volume if that volume is being formatted.
5. You cannot resize (shrink) a volume that is being migrated.
6. You can not resize (shrink) a volume if cloud snapshot is enabled on that volume.

This command is supported for volumes in Metro Mirror and Global Mirror relationships that are in `consistent_synchronized` state if those volumes are using thin-provisioned or compressed copies. This command is not supported for volumes:

- In HyperSwap relationships or in Global Mirror relationships that are operating in `cycling` mode
- In relationships where that volume has a Change Volume configured
- That have a fully-allocated copy

You must shrink both volumes in a relationship to maintain full operation of the system. To perform this:

1. Shrink the secondary volume by the required capacity
2. Shrink the primary volume by the required capacity

An invocation example to decrease the capacity of vdisk1 by 2 KB

```
shrinkvdiskspace -size 2048 -unit b vdisk1
```

The resulting output:

No feedback

An invocation example to decrease the capacity of vdisk2 by 100 MB

```
shrinkvdiskspace -size 100 -unit mb vdisk2
```

The resulting output:

No feedback

An invocation example to decrease the real capacity of thin-provisioned vdisk3 by 100 MB without changing its virtual capacity

```
shrinkvdiskspace -rsize 100 -unit mb vdisk3
```

The resulting output:

No feedback

An invocation example to decrease the real capacity of thin-provisioned VDisk copy ID 1 of mirrored vdisk3 by 100 MB

```
shrinkvdiskspace -rsize 100 -unit mb -copy 1 vdisk4
```

The resulting output:

No feedback

An invocation example to decrease the virtual capacity of thin-provisioned vdisk5 by 1 GB without changing its real capacity

```
shrinkvdiskspace -size 1 -unit gb vdisk5
```

The resulting output:

No feedback

rmvolume

Use the **rmvolume** command to remove a volume. You can use this command for high availability configurations that include HyperSwap or stretched systems.

Syntax

```
➤➤ rmvolume — [ -removehostmappings ] [ -removercrelationships ] [ -removefcmaps ]
➤ [ -discardimage ] [ -cancelbackup ] [ volume_id ] [ volume_name ] ➤➤
```

Parameters

-removehostmappings

(Optional) Allows a volume to be deleted even if host mappings are removed if this volume is deleted.

-removercrelationships

(Optional) Allows a volume to be deleted even if it is part of a remote copy relationship.

-removefcmaps

(Optional) Allows a volume to be deleted even if it is part of a FlashCopy mapping and regardless of the state of the mappings. FlashCopy mappings that are `rc_controlled` (for change volumes) require this parameter to be specified to force the deletion of a change volume while it is configured in a remote copy relationship. However, it is recommended to remove the change volume from the relationship before you delete it to avoid data loss. HyperSwap volumes with only the `rc_controlled` FlashCopy mappings for change volumes do not require this parameter to be specified.

-discardimage

(Optional) Allows a volume to be deleted even if data on an image mode copy cannot be made consistent.

Important: Using this parameter might result in data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of data loss on the volume.

-cancelbackup

(Optional) Allows a volume to be deleted even if a backup operation is in progress.

Important: Using this parameter might result in data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of data loss on the volume.

volume_id | volume_name

(Required) Specifies the volume to be removed.

Description

Use the **rmvolume** command to remove a volume.

For a HyperSwap volume, the active-active relationship and the change volumes are also deleted.

An invocation example to remove a volume

```
rmvolume 0
```

The detailed resulting output:

No feedback

An invocation example to remove a volume with FlashCopy mappings

```
rmvolume -removefcmaps 1
```

The detailed resulting output:

No feedback

An invocation example to remove a master or auxiliary volume in a Global Mirror relationship with change volumes

```
rmvolume -removercrelationships 6
```

The detailed resulting output:

No feedback

An invocation example to remove a HyperSwap volume

```
rmvolume myhyperswapvol
```

The detailed resulting output:

No feedback

An invocation example

```
rmvolume -cancelbackup 1
```

The detailed resulting output:

No feedback

An invocation example

```
rmvolume -discardimage 1
```

The detailed resulting output:


No feedback

rmvolumecopy


Use the **rmvolumecopy** command to remove a volume copy from a volume. You can use this command for high availability configurations that include HyperSwap or stretched systems.

Syntax

This syntax diagram specifies a volume copy by site.



This syntax diagram specifies a volume copy by copy ID or storage pool. You must specify either **-copy** or **-pool**.



Parameters

-site *site_id* | *site_name*

(Required) Specifies the site that the volume copy is removed from. You cannot specify this parameter if you specify **-copy** or **-pool**.

-pool *storage_pool_id* | *storage_pool_name*

(Optional) Specifies the storage pool that the volume copy is removed from.

-copy *copy_id*

(Optional) Specifies the copy ID for the volume copy to be deleted. The value is 0 or 1. This keyword cannot be specified whether **-site** is specified.

Important: If the volume has copies in multiple sites you must specify **-pool** to identify the volume copy to remove.

-removefcmaps

(Optional) Allows a volume copy to be deleted even if it is part of a FlashCopy mapping.

Important: Use this parameter for HyperSwap volumes.

-discardimage

(Optional) Allows a volume copy to be deleted even if data on an image mode copy cannot be made consistent.

Important: Using this parameter might result in data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of data loss on the volume.

-force

Deletes the specified volume copy whether a cloud system backup or restore operation is in progress. If you do not specify this parameter a volume or volume copy cannot be deleted while a backup operation is in progress. Additionally, a volume containing image mode copies cannot be deleted while a restore operation is in progress (if the volume contains inconsistent data).

Important: Using this parameter might result in data loss. Use it only under the direction of your product support information, or if you are willing to accept the risk of data loss on the volume.

volume_id | volume_name

(Required) Specifies the volume ID or name for the volume copy to remove. The value for volume ID must be a number and the value for volume name must be an alphanumeric string.

Description

Use the **rmvolumecopy** command to remove a copy of a volume.

For a HyperSwap volume, the active-active relationship and the change volumes are also deleted.

Remember: HyperSwap volumes that are part of a consistency group must be removed from that consistency group before you can remove the last volume copy from that site.

This command fails if a backup operation is in progress.

An invocation example to remove a volume copy at site 1 on a HyperSwap system

```
rmvolumecopy -site 1 0
```

The detailed resulting output:

No feedback

An invocation example to remove a volume copy when there are two copies in the same storage pool.

```
rmvolumecopy -pool 5 -copy 1 volume5
```

The detailed resulting output:

No feedback

An invocation example to remove a volume copy with FlashCopy mappings.

```
rmvolumecopy -site 1 -removefcmaps 1
```

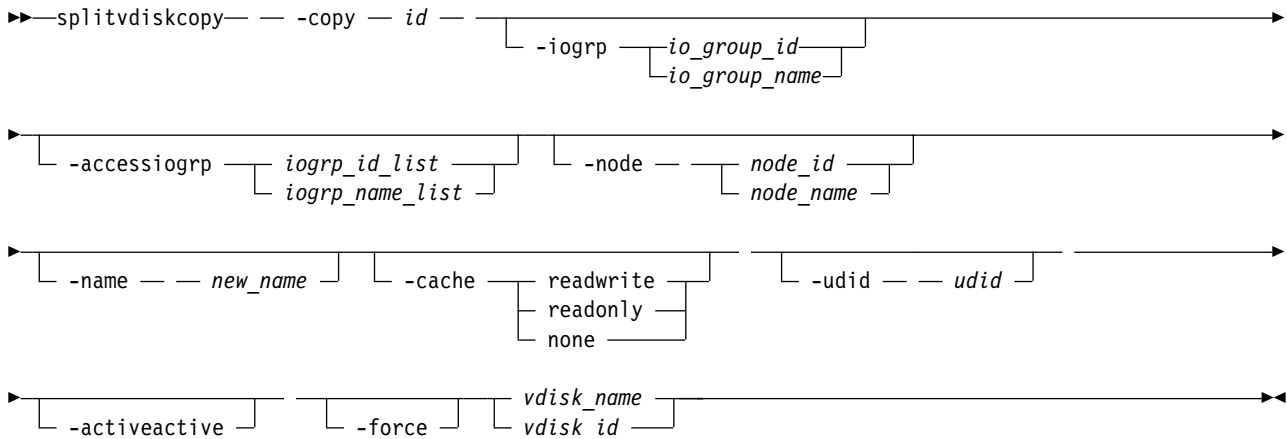
The detailed resulting output:

No feedback

splitvdiskcopy

Use the **splitvdiskcopy** command to create a separate volume from a synchronized copy of a mirrored volume.

Syntax



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 volume to. The default is the I/O group of the specified volume.

-accessiogrp *iogroup_id_list* | *iogroup_name_list*

(Optional) Specifies which I/O groups provide access to the volume. If the **-accessiogrp** parameter is used, the specified I/O groups provide access even if that set includes either the caching I/O group of the original volume or the caching I/O group of the new volume. If the flag is not specified and the original volume has only its caching I/O group in the set of I/O groups that provide access to the original volume, the new volume is assigned its caching I/O group as the only I/O group that provides access (which might not be the same as caching I/O group of the original volume).

Otherwise, the new volume provides access using the same set of I/O groups that are used with the original mirrored volume.

Note: The I/O groups that are specified are not required to include the caching I/O group.

-node *node_id* | *node_name*

(Optional) Specifies the preferred node ID or the name for I/O operations to this volume. You can use the **-node** parameter to specify the preferred access node.

-name *new_name*

(Optional) Assigns a name to the new volume.

-cache *readwrite* | *readonly* | *none*

(Optional) Specifies the caching options for the new volume. (Optional) Specifies the caching options for the volume. Valid entries are:

- *readwrite* to enable the cache for the volume
- *readonly* to disable write caching while allowing read caching for a volume
- *none* to disable the cache mode for the volume

The default is *readwrite*.

Remember: If you do not specify the **-cache** parameter, the default value (*readwrite*) is used.

-udid *udid*

(Optional) Specifies the *udid* for the new volume. The *udid* is a required identifier for OpenVMS

hosts; no other hosts use this parameter. Supported values are a decimal number 0 - 32767, or a hexadecimal number 0 - 0x7FFF. A hexadecimal number must be preceded by 0x; for example, 0x1234. The default *udid* value is 0.

-activeactive

(Optional) Specifies that an active-active relationship is created between the specified volume and the newly created volume.

-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 volume might not be consistent.

Description

The **splitvdiskcopy** command creates a new volume in the specified I/O Group from a copy of the specified volume.

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 command failure, wait for the copy to synchronize or split the unsynchronized copy from the volume by using the **-force** parameter. You can run the command when either volume copy is offline.

For active-active relationships, the existing volume must be in an I/O group with a site name or site ID. The existing volume must also use a storage pool with the same site information. The new volume must be created in an I/O group with a site name or site ID (that is not the same site as the I/O group for the existing volume). The new volume must use a storage pool with the same site name or site ID as the I/O group's site name or site ID. The topology must be hyperswap or the active-active relationship is not allowed.

For active-active relationships, the existing volume must not be the target of a FlashCopy mapping.

You can use this command to partially create a HyperSwap volume and:

1. Configure the access I/O groups for the existing volume to include the new volume's I/O group
2. Create and associate change volumes to the active-active relationship

When these tasks are completed, the active-active relationship can start or resynchronize any regions that are written to on the existing volume. The created relationship uses the existing volume as its master copy, and the new volume as its auxiliary copy.

Note: If the I/O group has enough bitmap space available to allocate for remote copy and the allocated space for the remote copy is not large enough to accommodate the new relationship, space is automatically added. (Remote copy includes Global Mirror, Metro Mirror, and active-active relationships.)

An invocation example for creating a volume with I/O groups 2 and 3 in its I/O group access set

```
splitvdiskcopy -copy 1 -iogrp 2 -node 7 -accessiogrp 2:3 DB_Disk
```

The resulting output:

Virtual Disk, copy [1], successfully created.

An invocation example

```
splitvdiskcopy -copy 1 vdisk8
```

The resulting output:

Virtual Disk, id [1], successfully created.

An invocation example

```
splitvdiskcopy -activeactive -iogrp siteB -copy 1 -name siteBvolume siteAvolume
```

The resulting output:

Virtual Disk, copy [1], successfully created.

Chapter 30. Command-line interface messages

This section lists the messages that can be displayed while you use the command-line interface (CLI).

The CLI displays a return value upon completion of the command. If the command completes normally and without error, the return code is 0. If the command fails, the return code is **1** and the Error Code is sent to standard error. If the command succeeds, but the cluster is operating near its licensed virtualization limit, the return code can still be 1, and a warning Error Code is sent to standard error.

When a create command is issued, the message ID that has been assigned to the new object is returned as part of the success message sent to standard output. If the **-quiet** parameter is used, only the message ID is sent to standard output.

CMMVC4041E You must use an 0x parameter for the -lba and -vdisklba parameters.

Explanation: The parameter format for lsmdisklba needs explaining if using the incorrect format.

User response: Retry the command using the described parameter format.

CMMVC4042E You must use an 0x parameter for the -lba and -mdisklba parameters.

Explanation: The parameter format for lsmdisklba needs explaining if using the incorrect format.

User response: Retry the command using the described parameter format.

CMMVC4043E You must use an 0x parameter for the -mdisklba parameter.

Explanation: The parameter format for lsmdisklba needs explaining if using the incorrect format.

User response: Retry the command using the described parameter format.

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 (for example, the specified string is an incorrect path).

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 time stamp format is YYMMDDHHMMSS.

Explanation: The specified value is not a valid time stamp format. The valid format is YYMMDDHHMMSS.

User response: Use the correct time stamp format, and resubmit the command.

CMMVC5722E *VALUE* contains a month value that is not valid. The valid time stamp format is YYMMDDHHMMSS.

Explanation: The month value (MM) that you have specified is not valid.

User response: Specify a valid month value, and resubmit the command.

CMMVC5723E *VALUE* contains a day value that is not valid. The valid time stamp format is YYMMDDHHMMSS.

Explanation: The day value (DD) that you have specified is not valid.

User response: Specify a valid day value, and resubmit the command.

CMMVC5724E *VALUE* contains an hour value that is not valid. The valid time stamp format is YYMMDDHHMMSS.

Explanation: The hour value (HH) that you have specified is not valid.

User response: Specify a valid hour value, and resubmit the command.

CMMVC5725E *VALUE* contains a minutes value that is not valid. The valid time stamp format is YYMMDDHHMMSS.

Explanation: The minutes value (MM) that you have specified is not valid.

User response: Specify a valid minutes value, and resubmit the command.

CMMVC5726E *VALUE* contains a seconds value that is not valid. The valid time stamp format is YYMMDDHHMMSS.

Explanation: The seconds value (SS) that you have specified is not valid.

User response: Specify a valid seconds value, 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 %1 is not a valid time format. The valid time format is MMDDHHmmYYYY with YYYY<2070.

Explanation: The specified value should be in the format MMDDHHmmYYYY with YYYY less than 2070.

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 command that you have specified must be run on the configuration node.

User response: Log off of the node service IP address, log on to the management 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: Reduce the size of the file system by deleting obsolete logs or other unnecessary files, then retry the command.

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.

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 an update is in progress.

Explanation: The operation failed because an update is in progress.

User response: Wait for the update 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 recovery group does not support the command line interface.

User response: Use the management GUI or the **lsiogrp** command to see the list of active I/O groups. Rerun the command with a valid I/O group ID or name.

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 an update is in progress.

Explanation: The node was not deleted because an update is in progress.

User response: Wait for the update 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 update of the cluster 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 update of the cluster 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 update of the cluster could not proceed because there is an I/O group in the cluster that contains only one node. The update 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 update.

Explanation: The update of the cluster could not proceed because there is an I/O group in the cluster that contains only one node. The update 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 update.

User response: Either update 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 action failed because the specified resource was unavailable.

Explanation: The resource specified in the action was not available for use.

User response: Fix any errors associated with the specified resource, or reissue the command using an alternative resource.

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 action failed because the managed disk (MDisk) is not in managed mode.

Explanation: The action is only permitted on MDisk that are currently in the managed mode.

User response: Either add the MDisk to a storage pool, or specify a different MDisk.

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 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 volume was not deleted because it is mapped to a host, or it is part of a FlashCopy mapping or Remote Copy relationship, or a cloud snapshot or restore operation is in progress, or it is involved in an image mode migration.

Explanation: The volume was not deleted for one of the following reasons:

- It is mapped to a host
- It is part of a FlashCopy mapping
- It is in a remote copy relationship
- A cloud snapshot operation is in progress
- A restore operation is in progress

User response: Make sure that you specified the correct volume. If so, take one or more of the following actions:

- If the volume is mapped to a host, remove the host mapping.
- If the volume is part of a FlashCopy mapping, remove the mapping.
- If the volume is in a remote copy relationship, remove it from the relationship.
- If a cloud snapshot or restore operation is in progress, either wait for the operation to finish or cancel the operation.

When the volume is available for deletion, retry the delete 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 `lsfreeextents <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 host object was not deleted because it is mapped to one or more volumes.

Explanation: There is at least one volume mapped to this host. Therefore the host object cannot be deleted without the loss of data.

User response: Verify that you are attempting to delete the correct host. If so, use the **lshostvdiskmap** command to locate the volumes that are mapped to this host.

If you are sure that these mappings are no longer required, remove them using the **rmvdiskhostmap** command, and retry the original **rmhost** 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 fix 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 fix 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 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 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 fix 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, and the partner cluster is running a downlevel software version.

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.

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 (VDisk) 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 fix 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: Increase the memory with the chiogrp command.

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: This error may occur because one or more of the source or target volumes for relationships within the consistency group is offline. It might also be caused if one or more of the source or target volumes are inaccessible because it is participating in a FlashCopy mapping that is prepared or incomplete.

User response: Perform the following steps:

- Ensure that at least one of the nodes in the I/O group of each of the source and target volumes is online.
- Ensure that the source and target volumes for each relationship in the consistency group are online.
- Correct any problems that may be holding the source or target volumes offline. For example, a thin-provisioned volume may be held offline if all available space is allocated.
- Ensure that any FlashCopy mappings that involve the source or target volumes for each relationship are complete.
- Fix all of the unfixed events in the event log.

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 startrelationship command is required when the relationship is not synchronized because consistency would be lost by starting the copy operation. Submitting the 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 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 might be due to the file system being full.

Explanation: dump file was not created. This might be due to the file system being full.

User response: Not applicable.

CMMVC5984E The dump file was not written to disk. This might be due to the file system being full.

Explanation: The dump file was not written to disk. This might 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 update package does not exist.

Explanation: The specified update package does not exist.

User response: Not applicable.

CMMVC5994E Error in verifying the signature of the update package.

Explanation: The system could not verify the signature of the update 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 An error prevented the unpacking of the update package.

Explanation: The system disk is too full to allow the update package to be unpacked.

User response: Use the cleardumps command with the parameter -prefix /home/admin/upgrade/ to clear unused files, then reboot the node before attempting to unpack the update package again.

CMMVC5996E The specific update package cannot be installed over the current version.

Explanation: The update package is not compatible with the current version or the system.

User response: Check the available update packages and find the correct update package for your current version and for your system. If the update package is correct for your system, check the version requirements for the package. You might have to update the current version to an intermediate version before you update to the latest version. (For example, if your current version is 1 and you are trying to update to version 3, you might need to update to version 2 before applying the version 3 update.)

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 allocate 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 update 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 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 update pre-install process failed.

Explanation: The update 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 update package. Download the package from the source location again as it might have been corrupted during a network transfer.

CMMVC6019E The update failed because a node pended.

Explanation: The update failed because a node pended as the update was in progress.

User response: Ensure that all nodes are online and available before restarting the update process.

CMMVC6020E The update failed because the system was unable to distribute the package to all of the nodes.

Explanation: The system could not complete the process of updating files. A full disk is a possible cause.

User response: Ensure that all nodes are online, and use the cleandumps command to clean the updates directory.

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 package cannot be applied to the current code level because it contains changes to the cluster state and there are remote cluster partnerships defined.

Explanation: The action failed because there is a connected remote cluster. The update cannot be applied because it would put the remote cluster at a different code level than the local cluster.

User response: Ensure that the cluster partnership is unconfigured before resubmitting the command. Ensure that you unconfigure the remote cluster and update the code on it before reconfiguring the cluster partnership.

CMMVC6029E All nodes must have identical code level before a concurrent code update can be performed.

Explanation: The concurrent update failed as two or more nodes were at differing code levels. All nodes must be at the same code level before a software update can be performed.

User response: Use the service assistant to bring all nodes to the same level before resubmitting the concurrent update.

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 update is in progress.

Explanation: The action failed because a software update is in progress.

User response: Wait for the software update 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.

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.

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: Reissue 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 an update. Depending on the nature of the errors, they might cause the update 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 an update. Depending on the nature of the errors, they might cause the update 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.

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: Contact your administrator. Resubmit the command after the administrator confirms that the problem is resolved.

CMMVC6092E The metadata recovery operation was aborted.

Explanation: The metadata recovery operation was cancelled.

User response: Check your command and ensure that your input is correct. Resubmit the command. If you get the same error, contact your administrator.

CMMVC6093E Metadata recovery internal error - (read only)

Explanation: Metadata recovery internal error - (read only)

User response: Contact your administrator. When the administrator lets you know that the problem has been resolved, resubmit the command.

CMMVC6095E Metadata recovery encountered the end of the disk.

Explanation: Metadata recovery encountered the end of the disk.

User response: Contact the administrator and let them know about this error. The administrator must take care of this problem before you can continue.

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: Check your command. Correct the specified node and resubmit..

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 Backup file version version_id is not compatible with current version version_id

Explanation: The backup file was generated on a cluster with a different version number than that of your current system.

User response: Make sure that you are copying the most up-to-date backup file onto the cluster. If you believe your backup file is correct, contact your service support representative.

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. Specific steps are dependent on what command was running and what return code was given.

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 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: Check your command. Specify the correct configuration file and resubmit the command.

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: The command you submitted cannot be processed at the moment. Follow the message and submit a different command first.

CMMVC6146E Problem parsing *OBJECT_TYPE* data: *LINE*

Explanation: An unexpected error has occurred.

User response: Contact the support center.

CMMVC6147W *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.

CMMVC6157E Object is required.

Explanation: An object or target was not specified for the command.

User response: Refer to the documentation for the command and specify all required parameters.

CMMVC6164E *variable_error_message*

Explanation: This message is generated by the system and varies depending on the circumstances in which it was created.

User response: For help with a specific message, contact your service support representative.

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: Re-create the default cluster with the correct configuration node, and resubmit the command.

CMMVC6166E The property *PROPERTY* of the object *OBJECT* has changed during *svconfig restore -execute*.

Explanation: The integrity of the restoration cannot be guaranteed.

User response: Resubmit the command from *svconfig restore -prepare*.

CMMVC6169E *variable_error_message*

Explanation: This message is generated by the system and varies depending on the circumstances in which it was created.

User response: For help with a specific message, contact your service support representative.

CMMVC6171I *percentage%* metadata compression

Explanation: The message displays the extent of metadata compression.

User response: This message is informational only. No action is required.

CMMVC6174I Pausing for *num_minutes* minute(s) after node addition(s)

Explanation: The system must process the addition of one or more nodes for the amount of time that is displayed.

User response: This message is informational only. No action is required.

CMMVC6175I Resuming after pause

Explanation: The system completed processing the addition of one or more nodes.

User response: This message is informational only. No action is required.

CMMVC6180E *object_type* [*object_id*] not restored because [*object_type* | *object_property*] is missing

Explanation: The error message has two variants. One specifies a missing object type, such as a drive, and the other specifies a missing object property, such as a name. In either case, the system cannot continue with

object creation because of the missing object or property.

User response: Add the specified object or property to the system. If you believe that the object or property already exists, contact your service support representative.

CMMVC6180W *object_type object_identifier* not restored because object with property *property_value* is missing

Explanation: The system cannot continue with object creation because a dependent object is missing.

User response: Add the specified dependent object to the system. If you believe that the object already exists, contact your service support representative.

CMMVC6181E *object_type object_id* has property *property_value*; should be *property_value_in_backup* [; use **-force to override** | **overridden**]

Explanation: The displayed object has a property value in its configuration that does not match the value for that property in the backup.

User response: Check the properties in the current configuration. The error message might include one of the following additions:

use -force to override

If this version of the message is displayed, you can retry the command with the addition of the **-force** parameter to retain the property value that is in the current configuration.

overridden

This version of the message is displayed for informational purposes only. The property value in the current configuration is retained. No user action is required.

Otherwise, you must correct the property value in the current configuration before you retry the command.

CMMVC6182E *object_type object_name* with property *property_value* cannot be [restored | backed up].

Explanation: Either the displayed object was not restored, or was not backed up, as shown.

User response: The user response is different for each object type. Contact your service support representative.

CMMVC6186E *io_grp io_group_name* restored with id *new_id* instead of *old_id*.

Explanation: This situation can occur when the configuration node is different from the node that was used to create the original cluster. This change affects the SCSI Inquiry value for the I/O group.

User response: This message is a warning only. No user response is required.

CMMVC6187W Extraneous *object_type object_name* discovered

Explanation: Recovery identified an extraneous object of the displayed type and named it *object_name*.

User response: This message is only a warning. No action is required.

CMMVC6188W VARIANT #1: Current Config inconsistent with backup *object_type object_id* does not exist in backup
 VARIANT #2: Current Config inconsistent with backup *object_type object_id* now has wrong_property
wrong_value was correct_value

Explanation: VARIANT #1: The object of the displayed type and ID is in the current configuration but not in the backup file. The object will not be recovered.

VARIANT #2: The recovery process found an inconsistency. The object of the displayed type and ID now has a property with an incorrect value.

User response: VARIANT #1: Re-create the missing object in the backup file after the recovery process is complete.

VARIANT #2: Manually restore the correct value after the recovery process is complete.

CMMVC6189W E-mail server settings could not be restored due to missing configuration information

Explanation: An attempt was made to restore email server settings that were not found on the backup device.

User response: Make sure that you specified the correct email server and that you made no typographical errors. If everything looks correct, contact your service support representative.

CMMVC6200E Discovery did not complete within time limit - check that this is expected

Explanation:

User response: Contact IBM Support.

CMMVC6201E System layer must be changed to *new_layer* using chsystem before configuration can be restored

Explanation:

User response: Contact IBM Support.

CMMVC6202E This command can only be run by the superuser

Explanation:

User response: Contact IBM Support.

CMMVC6203E Enclosure serial number *serial_no* was found for enclosure ID *enclosure_1* but it is already in use for enclosure ID *enclosure_2*

Explanation:

User response: Contact IBM Support.

CMMVC6204E Node *node_name* in backup is part of *io_grp io_group_1*, but the node is currently set to *io_grp io_group_2*

Explanation:

User response: Contact IBM Support.

CMMVC6205E Cannot retrieve VPD for node *node_name* because it is currently in service mode.

Explanation:

User response: Contact IBM Support.

CMMVC6206E Timeout exceeded waiting for *object_type object_id* to have expected *property_name expected_value*, actual *actual_value*.

Explanation:

User response: Contact IBM Support.

CMMVC6207E Cannot restore encryption as more USB devices are required. Required *req_number_devices* device(s) but found *actual_number_devices* device(s).

Explanation:

User response: Add the required number of USB devices and retry the command.

CMMVC6208E The *feature_id* feature cannot be partially configured prior to a restore operation. The current state is *current_state*.

Explanation:

User response: Contact IBM Support.

CMMVC6209E Unable to determine system code version.

Explanation:

User response: Contact IBM Support.

CMMVC6210E Current Config inconsistent with backup because metadatavdisk with *vdisk_id volume_id* already exists.

Explanation:

User response: Contact IBM Support.

CMMVC6211E The system contains fewer USB ports than required to recover encryption automatically. Manually set up encryption and then run the **-prepare** command again.

Explanation:

User response: Manually set up encryption and then run the **-prepare** command again.

CMMVC6212E The host port mode did not change to transitional within the time limit.

Explanation: The value of the *fctargetportmode* variable should change to transitional within 1 minute. If it does not, the T4 recovery fails.

User response: Contact IBM Support.

CMMVC6213E Active-active remote copy relationship *relationship_id* not added to consistency group *group_id*.

Explanation:

User response: Contact IBM Support.

CMMVC6215W SVC client certificate must be exported and installed on all key servers before running execute.

Explanation: A new cluster was created as part of the T4 store process. The new cluster has a new cluster (client) certificate. This certificate must be exported and installed on all key servers so that subsequent recovery commands in the execute phase, such as **mkkeyserver**, will succeed.

User response: Be sure to export the new cluster

certificate by using the **chsystemcert -export** command and install the SSL certificate on all key servers before you proceed with the **svconfig restore -execute** command.

CMMVC6216E Failed to generate key server certificates from the backup file.

Explanation: As part of the T4 restore script, if a key server configuration existed in the backup, key server SSL certificates are regenerated and written to new files, which are then used by the restore process. This error means that the restore process failed to regenerate the key server certificates.

User response: Contact your service support representative to manually restore the key server configuration, including server certificates.

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 operation 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 operation 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 operation 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 operation 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: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: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: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

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 or compressed.

Explanation: You are attempting to run a command that is valid only for thin-provisioned or compressed volumes.

User response: Specify a thin-provisioned or compressed 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 virtual disk (VDisk) copy is already being repaired.

Explanation: You are attempting to repair a thin-provisioned or compressed 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 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 or compressed volume copy that is reporting corrupt metadata. The cluster cannot repair the volume copy because it was not recognized as a valid thin-provisioned or compressed 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 update 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 or formatting in the I/O group. Mirroring bitmaps are temporarily used for the formatting of non-mirrored volumes as well as tracking the synchronization of mirrored volumes.

User response: Perform one of the following actions:

- Submit a 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 a 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 a `rmvdiskcopy` or `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 `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 `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 a `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.

Explanation: An attempt was made to remove an image mode copy, and the data on the copy is not synchronized with the host-accessible copy.

User response: Wait for the volume copy to resynchronize, then retry the command. Alternatively, specify the **-discardimage** or the **-force** parameter to force the delete operation.

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 `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 a `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 `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 a `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 `lsvdisksyncprogress` 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 a `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 `lsvdisk` 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 code.

Explanation: The version of code 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 code that you want to install. update hardware so that all of the hardware in the cluster is supported by the new version of code, 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:

- update 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 update 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 code might be incompatible with the remote cluster.

Explanation: The 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: update 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: update 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: update 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: update 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: update 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: update 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 fix 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 MDisk 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.

CMMVC6421E The specified size exceeds the maximum permitted for this feature.

Explanation:

- A **chiogrp** command was attempted where the maximum bitmap memory for one or more functions was exceeded. The maximum supported amount of bitmap memory for the RAID function, the Volume Mirroring function and the Metro Mirror or Global Mirror Copy Services functions is 512 MB. The maximum supported amount of bitmap memory for the FlashCopy® function is 2048 MB.
- Alternatively, a remote copy was attempted that used a memory size of 512 MB when the value of the **remote_copy_free_memory** variable reached 0. This scenario is unlikely.

User response: Retry the command, specifying a smaller amount of memory where necessary.

CMMVC6422E The specified size is too big. Total size of memory across all features exceeds maximum permitted.

Explanation:

- The maximum combined amount of memory across all functions other than FlashCopy® is 552 MB. A **chiogrp** command was used to change the memory allocations for an I/O group where the specified **-size** value (which defaults to units of megabytes) caused the total memory to exceed the maximum.
- Alternatively, an attempt was made to create a large HyperSwap volume, but enough memory was not available to complete the request.

User response: Review the amount of memory that is used by each feature by running an **lsiogrp** command that specifies the name or ID of an I/O group. Run one or more **chiogrp** commands to redistribute the amount of memory that each feature uses without exceeding the maximum. For more information, run the **help chiogrp** 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: Check the object specified in the command and determine if you need to specify a different object. Make the correction and resubmit the command.

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 `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 `mkrcrelationship` 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 `mkrcrelationship` 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 MDisk. 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 MDisk. 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 code version that does not support this configuration.

Explanation: One scenario in which this error occurs is when a cluster at a higher version is partnered to a cluster at a lower version that does not support multiple cluster mirroring, and you attempt to create another partnership to a cluster at the higher version to implement multiple cluster mirroring. Adding a partnership to a third cluster is not supported while at least one cluster in the current partnership is at the lower version.

User response: Either update 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 code version that does not support this configuration.

Explanation: The code versions of the clusters in the existing partnership do not support a partnership with a cluster that has the code 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 version 5.1.0 or later that is already in partnership with another cluster at version 5.1.0 or later.

User response: Either update the downlevel cluster code 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 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 an SSH key and password were not specified for the remote authentication service.

Explanation: An SSH key and password are required for the remote authentication service.

User response: Specify a valid SSH key and password when issuing 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: Recheck your command and make the correction to the source or destination or both that you have specified. Then resubmit the command after you have made the correction.

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: Recheck your command for the specified I/O group and verify that it is correct. Make the correction and resubmit the command. If your original command is correct, do further research to correct this problem.

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: Recheck your command and ensure that you specified the correct consistency group. Make the correction and resubmit the command. If you specified the correct consistency group, then more research is necessary to correct this problem.

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: Recheck your command to ensure that you specified the FlashCopy mapping or consistency group that you intended. Make the correction and resubmit the command. If the command was correct, more research is needed before you can resubmit the command.

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.

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 `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.

CMMVC6510E The task has failed because the user name or password is not correct.

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 user name or 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 management 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 management 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.

Some other considerations to fix this error:

- Each distributed array occupies 16 slots, starting at an mdisk id divisible by 16, therefore you might also consider the need to delete 16 unwanted mdisks starting on a mdiskid boundary divisible by 16. If you need additional information see lsmdisk.
- Or, you can delete an unwanted distributed array to make space for a distributed array.
- You must also issue the detectmdisk command to ensure that the mdisk inventory is updated after deleting mdisks.

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 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 incorrect 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 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 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 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 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 software type is not correct, or the image file has been corrupted.

User response: Reinstall the software download image, and resubmit the task. If the problem persists, contact IBM technical support for assistance.

CMMVC6546E A drive error was detected during the Apply Drive Software task.

Explanation: At least one of the drives that was requested for update by the Apply Drive Software task has failed.

User response: Either remove the failed drives from

the list of drives for the Apply Drive Software task and resubmit the task, or perform problem determination on the failed drives.

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 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 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 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 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 lsnodetendvdisks 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 lsnodetendvdisks command for every node in the cluster; for all other commands you must submit the lsnodetendvdisks 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 MDisk and complete any outstanding fix procedures. 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 not online.

Explanation: The 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 -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 cannot be initiated because no download images were found in the package file.

Explanation: The drive software update package file was unpacked but no download software images were found in the package.

User response: Acquire a valid flash drive software update package file, and resubmit the task using the new package file.

CMMVC6568E The Apply Drive Software task cannot be initiated because for at least one of the specified drives, the specified file does not contain an image for that drive's technology.

Explanation: The package file documentation lists the drive types for which there are images.

User response: Acquire a valid drive software update 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 flash drive software update 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 ' () * + ,
- 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 unconfigure the remote authentication service while the service is enabled.

Explanation: You cannot unconfigure 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 `lsarraymembergoals` 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 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 cluster email system settings must be configured before you can submit a command for error notifications.

User response: Configure the cluster 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 update cannot start because one or more I/O groups are in maintenance mode.

Explanation: Maintenance mode is used during system servicing, which prevents updates.

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 nodes from this enclosure cannot be added to the specified I/O group, or another enclosure is in the process of being added.

Explanation: The nodes in the enclosure being added are used elsewhere in the cluster, the target I/O group contains nodes of a different control enclosure, or another enclosure has not yet completed the process of being added.

User response: If a node already exists in the I/O group, add nodes from the same enclosure only. If the I/O group is empty, you can use a different control enclosure in which the nodes are not in a cluster. If you are currently adding another enclosure, wait for the completion of that process. Ensure that both nodes of an added enclosure are online and that the enclosure is listed in the output of the `lsenclosure` 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 `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.

CMMVC6953E The action cannot be completed because volumes are dependent on the specified mdisk. Force is required.

Explanation: A volume is dependent on a drive that was specified on the `applydrivesoftware` command.

User response: If the drive is a member of a RAID0 array, consider whether to introduce additional redundancy to protect the data on that drive. If the drive is not a member of a RAID0 array, fix any errors in the event log that relate to the array. When the drive is a member of an array with sufficient redundancy, repeat the command. Alternatively, consider using the `-force` option.

Note: With any drive software update there is a risk that the drive might become unusable. Only use the `-force` option if you accept this risk.

CMMVC6972E The command cannot be initiated because the maximum number of extents for an MDisk would be exceeded.

Explanation: Each MDisk has a limited number of extents, which varies according to the extent size set by the `mdiskgrp`, this limit would be exceeded by this command.

User response: Create a different pool with a larger extent size by using `mkmdiskgrp`. Then, retry the command by using `mdiskgrp`.

Ensure that you are familiar with the maximum configurations of the SAN Volume Controller. Search for the term "configuration limits and restrictions" in the Search support and downloads search box at the www.ibm.com/support website. The relationship between extent size and the maximum MDisk capacity is shown in the Extents table.

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 maximum number of iSCSI qualified names (IQNs) plus WWPNS for the cluster is already configured.

Explanation: The command cannot be initiated because the maximum number of iSCSI qualified names (IQNs) plus WWPNS for the cluster has been reached.

User response: Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current iSCSI qualified name or WWPNS definitions are not required. Remove at least one iSCSI qualified name or WWPNS that is not required, and resubmit 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 `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 `lsmdisk` to list the MDisk, 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 ensure that all drives within the enclosure are online before resubmitting the command.

CMMVC7014E The command cannot be initiated because one or more of the drives are not supported for this RAID level.

Explanation: Only certain RAID levels are supported in some configurations.

User response: Consult the configuration guide to determine supported RAID levels.

CMMVC7015E The command cannot be initiated because one or more of the drives are located in the wrong node.

Explanation: For RAID 0, all of the members must be located in the same node. For RAID 1 or RAID 10, mirrored pairs must be located in different nodes.

User response: Consult the configuration guide to determine which drives to use for the selected RAID level.

CMMVC7016E Authorization has failed because the private key is not valid for the user name that you have specified.

Explanation: The private key and user name that you have provided do not match what has been defined on the cluster.

User response: Ensure that the private key is valid for the specified user name, and log in again.

CMMVC7017E Login has failed because the maximum number of concurrent CLI sessions has been reached.

Explanation: The cluster supports up to 10 concurrent CLI sessions. The login attempt would have exceeded the supported limit.

User response: Reduce the number of open CLI sessions, and log in again.

CMMVC7018E The command failed because the requested VDisk size is too large.

Explanation: The system has a maximum size for virtual disks (VDisks) that is currently 256 TB. While creating a new VDisk or resizing an existing VDisk, you requested a VDisk size that exceeds the maximum.

User response: Resubmit the command with a smaller VDisk size.

CMMVC7019E The command failed because the VDisk size is not a multiple of 512 bytes.

Explanation: VDisk capacity must be a complete number of blocks, where one block is 512 bytes. While creating a new VDisk or resizing an existing VDisk, you requested a VDisk size that is an incomplete number of blocks.

User response: Resubmit the command with a valid VDisk size.

CMMVC7020E The command failed because the maximum number of VDIs for this I/O group already exist.

Explanation: The system has a limit of VDIs per I/O group. A new VDisk cannot be created in an I/O group that has already reached the limit of VDIs.

User response: Choose another I/O group or delete some VDIs in this I/O group.

CMMVC7021E The command failed because the maximum number of VDisk copies already exist.

Explanation: The system has a limit on the number of VDisk copies that can be created. An additional VDisk copy cannot be created because the limit has been reached.

User response: Delete some existing VDisk copies and resubmit the command.

CMMVC7022E The command failed because NTP is active.

Explanation: You have attempted to manually set the cluster time while the cluster is configured to use NTP (network time protocol) to set its time.

User response: Disable NTP, and resubmit the command. If you are trying to set the time manually because the cluster time is incorrect, check the settings on the NTP server.

CMMVC7023E The command failed because the requested node name is in use as the failover name of another node.

Explanation: You have attempted to add a node to a cluster or rename a node that is already in the cluster. The new name that you have requested for the node is not valid because one of the nodes in the cluster has been configured with the requested new name as its failover name.

User response: Either resubmit the command specifying a different node name, or modify the configuration of the node in the cluster to change its matching failover name to a different failover name.

CMMVC7024E The command failed because the maximum number of file systems already exist.

Explanation: The maximum number of file systems has been reached. Additional file systems cannot be created.

User response: Remove an unused file system and reissue the command, or extend an existing file system by creating the VDisk there.

CMMVC7025E The command failed because the VDisk is associated with a file system and cannot be removed under your current user role.

Explanation: You are attempting to remove a VDisk that is associated with a file system. However, you do not possess the required role for file system actions and VDisk removal.

User response: Resubmit the task using the remove VDisk command.

CMMVC7026E The command failed because VDIsks exist in the file system.

Explanation: You are attempting to delete an MDisk group with which VDIsks are associated. The MDisk group cannot be removed while the associated VDIsks remain.

User response: Remove the file system VDIsks, and resubmit the command to remove the MDisk group.

CMMVC7027E The command failed because the requested action is not permitted on a VDisk that is in a file system.

Explanation: The VDisk that you have specified is associated with a file system, which disallows the requested action.

User response: The command cannot be completed on this VDisk. It will only succeed with a VDisk that is not associated with a file system.

CMMVC7028E The task cannot be completed because the FlashCopy target VDisk that you have specified is in a Metro Mirror or Global Mirror relationship, and the I/O group of the VDisk is different than that of the proposed FlashCopy mapping.

Explanation: The FlashCopy map must be in the same I/O group as the target VDisk because the VDisk is a component of a remote copy relationship.

User response: Specify the I/O group of the target VDisk when creating the FlashCopy map.

CMMVC7029E The task cannot be completed because one or more of the target VDIsks of the FlashCopy mappings is the primary of a mirroring Metro Mirror or Global Mirror relationship.

Explanation: The target VDisk is part of a remote copy relationship that is active.

User response: Either force stop the FlashCopy consistency group or stop any remote copy relationships.

CMMVC7030E The task cannot be completed because the target VDisk of the FlashCopy mapping is the primary of a mirroring Metro Mirror or Global Mirror relationship.

Explanation: The target of the FlashCopy map is a component of an active FlashCopy map.

User response: Either force stop the FlashCopy map or stop the remote copy relationship.

CMMVC7031E The task cannot be completed because the FlashCopy mapping target VDisk is a secondary in a Metro Mirror or Global Mirror relationship, or is the primary in an active relationship.

Explanation: The target VDisk of the FlashCopy map is part of an active remote copy relationship.

User response: Stop the remote copy relationship.

CMMVC7032E The task cannot be completed because one or more of the target VDIsks of the FlashCopy mappings is either a secondary in a Metro Mirror or Global Mirror relationship or the primary of an active relationship.

Explanation: A target VDisk of a FlashCopy map in the consistency group is part of an active remote copy relationship.

User response: Stop any remote relationships containing a target VDisk of a map in the consistency group.

CMMVC7033E The task failed because the current hardware configuration is not valid.

Explanation: You have issued the "chnodehw" command to enable new hardware that is faulty, unsupported, or incompletely installed.

User response: Follow service procedures as prompted by the management GUI to correct the hardware configuration. Then reissue the command.

CMMVC7036E The action failed because quorum is not permitted on the drive that you have specified.

Explanation: Quorum is only permitted on specific drive types. The drive that you have selected will not support quorum.

User response: Reissue the command specifying a different drive.

CMMVC7037E The action failed because the drive cannot be found.

Explanation: You have specified a drive does not appear to exist.

User response: Reissue the command specifying a different drive.

CMMVC7038E The action failed because the system was unable to initialize the quorum disk.

Explanation: A sequence of SCSI commands must be sent to a quorum disk before it can become available for use. One of these SCSI commands has failed.

User response: Fix any errors associated with the disk, choose a different resource for quorum, and then reissue the command.

CMMVC7039E The action failed because the specified drive is not online.

Explanation: The drive that you have specified is offline possibly as a result of errors.

User response: Fix any errors associated with the drive, or choose a different resource for quorum, then reissue the command.

CMMVC7040E The action failed because the specified MDisk is not online.

Explanation: The MDisk that you have specified is offline possibly as a result of errors.

User response: Fix any errors associated with the MDisk, or choose a different resource for quorum, then reissue the command.

CMMVC7041E The action failed because a better quorum candidate is available for use as quorum and override has not been enabled.

Explanation: Quorum disks are selected automatically based on a set of selection criteria. The resource that was selected is inferior to an alternative resource.

User response: Choose a different resource for quorum, or refer to the quorum documentation before using the -override parameter.

CMMVC7042E The action failed because the -override yes parameter was used without a specified drive or MDisk.

Explanation: The -override yes parameter must specify a drive or MDisk.

User response: Reissue the command with the correct syntax.

CMMVC7043E The action failed because the required extents could not be allocated.

Explanation: When an MDisk is specified for quorum, some extents must be allocated for use by the quorum disk. Sufficient extents were not available.

User response: Reissue the command using a different MDisk or migrate data from the MDisk to free up sufficient extents.

CMMVC7044E The action failed because the specified drive is either degraded or excluded.

Explanation: The drive that you have specified either contains errors or is in the Excluded state.

User response: Fix any errors associated with the drive, or choose a different resource for quorum, then reissue the command.

CMMVC7045E The action failed because the specified MDisk is either degraded or excluded.

Explanation: The MDisk that you have specified either contains errors or is in the Excluded state.

User response: Fix any errors associated with the MDisk, or choose a different resource for quorum, then reissue the command.

CMMVC7046E The action failed because the -rsize option must be set to auto.

Explanation: You have run the **mkvdisk** or **addvdiskcopy** command to import a compressed VDisk (using -compressed and -import). The **-rsize** option must be set with a value of **auto**.

User response: Resubmit the command with **-rsize auto**.

CMMVC7047E The action failed because the validate parameter is not supported for compressed VDIs.

Explanation: The command **repairsevdiskcopy -validate** was issued against a compressed VDisk. Unlike thin-provisioned VDIs, compressed VDIs do not support a validation function.

User response: None.

CMMVC7048E The action failed because the compressed VDisk copies are not all corrupted.

Explanation: You have issued the **repairsevdiskcopy** or **recovervdisk -copy** command against a compressed VDisk copy that is not marked as corrupt. Unlike thin-provisioned VDisk copies, the repair process for

compressed VDisk copies can only be run if the system detects that they are corrupted.

User response: The issued command is not required. If the VDisk is offline, consult the *Troubleshooting Guide* to resolve the problem.

CMMVC7049E The command failed because VDIsks are obstructing resources required by the compression function.

Explanation: Compression could not be enabled because a VDisk prevented internal resources from being reassigned from cache. A VDisk is offline or data could not be flushed from the cache quickly enough.

User response: If any VDIsks are offline, follow service procedures to bring them online before resubmitting the command.

CMMVC7050E The command failed because at least one node in the I/O group does not support compressed VDIsks.

Explanation: Compression is supported only on the CF8 and later models of San Volume Controller node. You have attempted to create a compressed VDisk in an I/O group that contains at least one node that does not meet these requirements.

User response: Resubmit the command with a different I/O group.

CMMVC7051E The command failed because the I/O group contains compressed VDIsks. The node being added does not support compressed VDIsks.

Explanation: Compression is supported only on the CF8 and later models of San Volume Controller node. You have attempted to add a node that does not support compression to an I/O group that already contains at least one compressed VDisk.

User response: Add the node to a different I/O group or add a different node to the specified I/O group.

CMMVC7052E The nested group search parameter is not valid for the target LDAP server type.

Explanation: The LDAP server type that you have specified is predefined to perform nested group search.

User response: Check your command to ensure that you have specified the correct type. Remember that the following rules apply to type and -nestedgroupsearch:

- If the type is **itds**, -nestedgroupsearch cannot be processed
- If the type is **ad**, -nestedgroupsearch can only be set to **client** or **off** because there is no server support.

- If the type is **other**, the -nestedgroupsearch parameter is fully configurable

After making your correction, resubmit the command.

CMMVC7053E The task cannot be initiated because the nested group search value (server) is not valid for the target LDAP server type.

Explanation: The LDAP server type that you have specified only supports client-side nested group search.

User response: Reissue the task specifying client-side nested group search.

CMMVC7054E The task cannot be initiated because the user name or password of the LDAP administrator was not specified.

Explanation: The user name and password of the LDAP administrator are not configured on the cluster as required. Once the credentials are configured, the user name and password can be changed separately.

User response: Reissue the task specifying both the LDAP administrator user name and password.

CMMVC7055E The task cannot be initiated because the specified IP address, port, and base distinguished name (DN) are already configured on an LDAP server.

Explanation: The same IP address, port, and base DN exist on more than one LDAP server.

User response: Reissue the task specifying an different IP address, port, and base DN.

CMMVC7056E The task cannot be initiated because the number of LDAP servers has reached the supported maximum.

Explanation: The cluster limits the number of LDAP servers that can be configured, and the limit has been reached. To remove a configured LDAP server, the rmldapserver command can be submitted.

User response: Remove a configured LDAP server, and resubmit the task.

CMMVC7057E The task cannot be initiated because the specified LDAP server is the only configured LDAP server.

Explanation: Removing the specified LDAP server would result in the failure of the remote authentication service.

User response: Disable the remote authentication service by submitting the chauthservice command, and reissue the task.

CMMVC7058E The task cannot be initiated because no LDAP server is configured.

Explanation: The LDAP remote authentication service cannot be used until at least one LDAP server has been configured. To configure LDAP servers, the `mkldapserver` command can be submitted.

User response: Configure a valid LDAP server, and reissue the task.

CMMVC7059E The task cannot be initiated because some remote users are not configured with an SSH key and password for the specified remote authentication service.

Explanation: An SSH key and password are required for all users of the remote authentication service. To identify remote users without an SSH key and password, the `lsuser` command can be submitted. To configure the user's authentication settings, you can use the `chuser` command.

User response: Configure the remote users with an SSH key and password or configure the users as local.

CMMVC7060E The task cannot be initiated because the parameters that you have specified are not valid for the LDAP authentication service.

Explanation: The authentication service URL, user name, password, and SSL certificate are not configurable for the LDAP authentication service.

User response: Reissue the task specifying valid parameters.

CMMVC7061E The task cannot be initiated because the LDAP administrator user name that you have specified is not valid.

Explanation: LDAP administrator user names must be a valid Distinguished Name, NT login, or User Principal Name.

- Distinguished Names must be a sequence of attribute=value pairs separated by a comma (,), semicolon (;), or plus sign (+), and include special characters and UTF-8 characters that are appropriately escaped with a backslash (\).
- NT logins are valid for Active Directory only and should be in the format `DOMAIN\user`. They must not start or end with a period (.) and both `DOMAIN` and `user` must exclude characters in the set: `\ / : ? " < > |`
- UPN logins are valid for Active Directory only and must be in the format `user@suffix`. Both `user` and `suffix` must exclude spaces and the following characters: `() < > , ; : \ " [] @`

User response: Reissue the task specifying a valid

Distinguished Name, NT login, or User Principal Name.

CMMVC7062E The task cannot be initiated because you have specified an LDAP attribute that is not valid.

Explanation: An LDAP attribute name can contain only alphanumeric characters and hyphens, and the name must begin with a letter.

User response: Reissue the task specifying a valid LDAP attribute name.

CMMVC7063E The task cannot be initiated because the Distinguished Name that you have specified is not valid.

Explanation: A Distinguished Name must be a sequence of attribute=value pairs separated by a comma (,), semicolon (;), or plus sign (+) that includes special characters and UTF-8 characters escaped with a backslash (\).

User response: Reissue the task specifying a valid Distinguished Name.

CMMVC7064E User authentication failed because one or more LDAP servers could not be contacted.

Explanation: The LDAP server is not operating correctly or an incorrect IP address and port are defined for the LDAP authentication service. The event log has been logged and the corresponding service procedure can be used to resolve this issue. To change the IP address and port of an LDAP server, a security administrator role can submit the `chldapserver` command.

User response: Ensure that the LDAP servers are operating correctly. Ensure that the IP address and port defined for each LDAP server is correct, and reissue the task.

CMMVC7065E User authentication failed because a timeout occurred while communicating with one or more LDAP servers.

Explanation: A timeout occurred while the cluster was attempting to contact the LDAP servers. This timeout might be caused by a TCP/IP network problem, the LDAP servers not operating correctly, or by an incorrect IP address and port being defined for the LDAP servers. The event has been logged and the corresponding service procedure can be used to resolve this issue. To change the IP address and port of an LDAP server, a security administrator can use the `chldapserver` command.

User response: Ensure that LDAP servers and the TCP/IP network between them and the cluster are functional. Ensure that the IP address and port defined

for each LDAP server is correct, and reissue the task.

CMMVC7066E User authentication failed because an SSL connection could not be established with one or more LDAP servers.

Explanation: An incorrect LDAP security configuration exists on the cluster or an SSL certificate on the cluster is not valid. The event was logged and the corresponding service procedure is available to resolve this issue. To turn off transport layer security, a security administrator can submit the `chldap` command or submit the `chldapserver` command to set the SSL certificate for an LDAP server.

User response: Ensure that the SSL configuration on each LDAP server is correct and that the SSL certificate defined in the cluster for each LDAP server is correct, or ensure that Transport Layer Security is disabled. Then reissue the task.

CMMVC7067E User authentication failed because one or more LDAP servers rejected an anonymous bind attempt.

Explanation: A user name and password were not specified on the cluster for LDAP authentication and the LDAP server has refused an attempt to bind anonymously. The event has been logged and the corresponding service procedure can be used to resolve this issue. To configure a user name and the password for LDAP authentication, a security administrator can submit the `chldap` command.

User response: Ensure that all LDAP servers are configured to allow anonymous bind, or configure a user name and password for LDAP authentication. Then reissue the task.

CMMVC7068E User authentication failed because one or more LDAP servers rejected an attempt to bind with the LDAP administrator credentials configured on the cluster.

Explanation: A user name and password were configured on the cluster for LDAP authentication and an LDAP server has refused an attempt to bind with these credentials. The event has been logged and the corresponding service procedure can be used to resolve this issue. To change the user name and password defined on the cluster, a security administrator can submit the `chldap` command.

User response: Ensure that the LDAP credentials configured on the cluster match the credentials that are configured on all LDAP servers, and reissue the task.

CMMVC7069E User authentication failed because one or more LDAP servers report an incorrect user name or password.

Explanation: The user name and password that you have provided do not match any user name and password on the configured LDAP servers. If the password for the user name has recently changed on the configured LDAP servers, it may be necessary to force the cluster to refresh its authentication cache. To force a refresh, a security administrator can submit the `chauthservice -refresh` command.

User response: Ensure that the user name and password are correct. Ensure that any recently changed passwords are flushed from the cache of the cluster, and reissue the task.

CMMVC7070E User authentication failed because the LDAP user attribute is incorrectly configured on one or more LDAP servers.

Explanation: The LDAP configuration on the cluster specifies an LDAP user attribute that does not exist on the LDAP server. Users cannot be identified by user name because the attribute is incorrectly configured. The event has been logged and the corresponding service procedure is available to resolve this issue. To specify a different user attribute, a security administrator can submit the `chldap` command.

User response: Ensure that the LDAP user attribute specified on the cluster is correct. Ensure that the schema on the configured LDAP servers includes the specified attribute, and reissue the task.

CMMVC7071E User authentication failed because the LDAP group attribute is incorrectly configured on one or more LDAP servers.

Explanation: The LDAP configuration on the cluster specifies an LDAP group attribute that does not exist on the LDAP server. The groups to which a user belongs cannot be identified because the attribute is incorrectly configured. The event has been logged and the corresponding service procedure can be used to resolve this issue. To specify a different group attribute, a security administrator can submit the `chldap` command.

User response: Ensure that the LDAP group attribute specified on the cluster is correct. Ensure that the schema on the configured LDAP servers includes the specified attribute, and reissue the task.

CMMVC7072E User authentication failed because the LDAP group attribute is not in a valid format on one or more LDAP servers.

Explanation: The LDAP group attribute in the user entry on the configured LDAP servers is in an invalid format. The event has been logged and the corresponding service procedure can be used to resolve this issue. The attribute must be a multivalued attribute containing the distinguished names of groups, or a colon-separated list of up to eight user group names.

User response: Ensure that the LDAP group attribute is correctly formatted on the LDAP servers, and reissue the task.

CMMVC7073E User authentication failed because the LDAP audit log attribute is not configured correctly on one or more LDAP servers.

Explanation: The LDAP configuration on the cluster specifies an LDAP audit log attribute that does not exist on the LDAP server. The string to use in the audit log cannot be identified because this attribute is incorrectly configured. The event has been logged and the corresponding service procedure can be used to resolve this issue. To specify a different audit log attribute, a security administrator can issue the chldap command.

User response: Ensure that the LDAP audit log attribute is correctly specified on the cluster. Ensure that the schema on the LDAP servers includes the specified attribute, and reissue the task.

CMMVC7074E The task cannot be initiated because the user could not be found on any of the configured LDAP servers.

Explanation: The remote user is configured but either no entry for the user exists on the configured LDAP servers, or more than one entry was found. The event has been logged and the corresponding service procedure can be used to resolve this issue.

User response: Ensure that the user name is unique on the LDAP servers. Ensure that the LDAP bind credentials allow the LDAP server to be searched, and reissue the task.

CMMVC7075I The LDAP task completed successfully.

Explanation: The LDAP task completed successfully.

User response: None.

CMMVC7076E *VOLUME* cannot be created with *VALUE* without *VALUE*.

Explanation: You are attempting to create a thin provisioned file system volume without compression. Thin provisioned file system volumes must include compression.

User response: Create a thin provisioned file system volume with compression or create a file system volume without thin provisioning.

CMMVC7077E The command failed because adding a thin provisioned copy to a file system volume is not allowed.

Explanation: You are attempting to add a volume copy to a file system volume that is not compressed but is thin provisioned. Only copies with compression or without thin provisioning can be added to file system volumes.

User response: Add a copy either with compression or without thin provisioning to the file system volume.

CMMVC7078E The command cannot be initiated because adding a copy to the storage pool of file system VDisks is not allowed.

Explanation: You are attempting to add a volume copy to a file system volume from a different storage pool. Only copies from the same storage pool can be added to file system volumes.

User response: Add a volume copy to a storage pool within the same file system volume, only.

CMMVC7079E The command failed because a volume copy must be different when added to a file system volume.

Explanation: You are only allowed to add a different volume copy to perform conversions between uncompressed and compressed.

User response: Add a compressed copy to a file system volume with an uncompressed copy or an uncompressed copy to a file system volume with a compressed copy.

CMMVC7080W The compressed storage used by the cluster is approaching the capacity that is licensed.

Explanation: You are being informed that compressed storage in use by the cluster has nearly reached the total licensed capacity.

User response: Compare your actual and planned usage of compression.

CMMVC7081W The compressed storage used by the cluster has exceeded the capacity that is licensed.

Explanation: You are being informed that compressed storage in use by the clustered system exceeds the total licensed capacity.

User response: Reduce the use of compressed storage or purchase additional licensing.

CMMVC7082W The number of control enclosures with compressed VDisks exceeds the number that are licensed.

Explanation: You are being informed that the licensed number of control enclosures that can contain compressed VDisks has been exceeded.

User response: Reduce or consolidate compressed VDisks or purchase additional licensing.

CMMVC7083E The specified number of control enclosures is not valid.

Explanation: The valid range of values for licensed control enclosures is 0–4. The value that you specify must be within this range.

User response: Specify a value between 0 and 4.

CMMVC7084E The action failed because the command is not permitted for compressed VDisks.

Explanation: The command that you have submitted is not valid on compressed volumes.

User response: Do not submit this command for a compressed volume.

CMMVC7102E The action could not be performed because one or more of the requested secondary VDisk is the target of an active FlashCopy mapping.

Explanation: An auxiliary volume cannot be selected as a change volume of a Remote Copy relationship while currently defined for a different relationship.

User response: Choose a different auxiliary volume.

CMMVC7143E The command cannot be initiated because nodes from another cluster are visible.

Explanation: The system layer can be changed only when no other systems are visible on the fabric.

User response: Change the Fibre Channel SAN zoning to remove connectivity between the nodes in the local system and nodes in remote systems, issue 'svctask detectmdisk', then retry the command.

CMMVC7144E The command cannot be initiated because it is not supported by this hardware type.

Explanation: The system layer can only be changed on Storwize family systems.

User response: None.

CMMVC7145E The command cannot be initiated because one or more partnerships are defined.

Explanation: The system layer can only be changed when there are no partnerships defined with remote systems.

User response: Remove all partnerships to remote systems, first removing any associated remote copy relationships and consistency groups using those partnerships, then retry the command.

CMMVC7146E The command cannot be initiated because a host object is associated with SAN Volume Controller ports.

Explanation: The system layer cannot be changed when there are host objects that contain Fibre Channel ports from SAN Volume Controller nodes or Storwize family systems.

User response: Remove all host objects that contain node ports, then retry the command.

CMMVC7147E The command cannot be initiated because one or more MDisk are provided by a storage system that does not support the system layer changing.

Explanation: The system layer cannot be changed if MDisk are currently being provided by a Storwize system.

User response: For each MDisk provided by a Storwize system, remove that MDisk from its storage pool. Once all such MDisk have been removed, change the Fibre Channel SAN zoning to remove connectivity between the nodes in the local system, and the nodes in the remote Storwize system. Finally, execute the **detectmdisk** command, and then change the system layer.

CMMVC7154E The task cannot be completed because the specified FlashCopy mapping is controlled by a Metro Mirror or Global Mirror relationship.

Explanation: The Metro Mirror or Global Mirror relationship under which the specified FlashCopy mapping is controlled prevents completion of the task.

User response: Check whether the specified task is permitted under the configuration of the Metro Mirror

or Global Mirror relationship.

CMMVC7155E The Create FlashCopy mapping task cannot be initiated because the source or target VDisk is being used as a change VDisk for a Metro Mirror or Global Mirror relationship.

Explanation: A volume cannot become the source or target of a FlashCopy mapping while in use as a change volume in a Metro Mirror or Global Mirror relationship.

User response: Specify a volume other than the source or target that is currently in use.

CMMVC7156E The change VDisk could not be associated because it is already a source or target VDisk in an existing FlashCopy mapping.

Explanation: A change volume cannot be associated if the same volume is the source or target of a FlashCopy mapping.

User response: Specify a volume other than the source or target that is currently in use.

CMMVC7157E The change VDisk could not be associated because the Metro Mirror or Global Mirror relationship has a VDisk on this cluster that is the target of a FlashCopy mapping in a different I/O group.

Explanation: The I/O group of the change volume conflicts with an I/O group in the relationship with which an association was attempted.

User response: Ensure that no conflicting I/O groups exist.

CMMVC7158E The change VDisk could not be associated because the Metro Mirror or Global Mirror relationship has a VDisk on this cluster that is already involved in the maximum number of FlashCopy mappings.

Explanation: The change volume cannot cause another volume to exceed the number of FlashCopy mappings allowed.

User response: Reduce the number of FlashCopy mappings in the volume that has reached its maximum.

CMMVC7159E The change volume could not be associated because the relationship has a volume on this cluster that is in an I/O group with no online nodes or because there are unrecovered FlashCopy mappings in the I/O group.

Explanation: FlashCopy metadata was lost and cannot be recovered.

User response: The nodes of the I/O group must be brought online before the change volume can be associated. If the nodes of the I/O group were removed, then delete any FlashCopy mappings that existed prior to removing the nodes. .

CMMVC7160E The change VDisk could not be associated because the I/O group has insufficient free bitmap space.

Explanation: The I/O group must have additional bitmap space to allow the change volume to be associated.

User response: Increase the total bitmap space of the I/O group.

CMMVC7161E The change VDisk could not be associated because the master change VDisk can only be configured from the master cluster, and the auxiliary change VDisk from the auxiliary cluster. The change VDisk must be configured from the remote cluster.

Explanation: A change volume must be associated from a cluster of the same type (master or auxiliary).

User response: Configure the change volume from the remote cluster.

CMMVC7162E The change VDisk could not be associated because one is already configured for the specified Metro Mirror or Global Mirror relationship.

Explanation: A change volume has been previously configured for the specified Metro Mirror or Global Mirror relationship.

User response: Ensure that the change volume is associated where a change volume has not been configured.

CMMVC7163E The change VDisk could not be associated because it is already involved in a Metro Mirror or Global Mirror relationship.

Explanation: The change volume is currently associated with a Metro Mirror or Global Mirror relationship.

User response: Choose an unassociated change volume for the specified Metro Mirror or Global Mirror relationship.

CMMVC7164E The change VDisk could not be associated because its size is different from those in the Metro Mirror or Global Mirror relationship.

Explanation: A change volume cannot be associated with volumes of a different size.

User response: Choose a change volume with a size that matches those with which it is being associated.

CMMVC7165E The change VDisk could not be disassociated because the Metro Mirror or Global Mirror relationship does not have one configured.

Explanation: An attempt was made to disassociate a change volume where one does not currently exist.

User response: Verify whether the intended change volume was specified.

CMMVC7166E The change VDisk could not be disassociated because it is currently in use by the Metro Mirror or Global Mirror relationship.

Explanation: An attempt was made to disassociate a change volume that is currently in use.

User response: Verify whether the intended change volume was specified.

CMMVC7167E The change VDisk could not be associated because it is mapped to a host.

Explanation: A change volume cannot be associated if it is mapped to a host.

User response: Unmap the change volume from its host, or choose a different change volume.

CMMVC7168E The VDisk-to-host mapping was not created because the VDisk is a change VDisk for a Metro Mirror or Global Mirror relationship.

Explanation: A volume cannot be mapped to a host if it is associated as a change volume.

User response: Choose a different change volume.

CMMVC7169E The Remote Copy relationship could not be deleted because this would corrupt the secondary VDisk.

Explanation: The deletion of the relationship is being prevented as a safeguard against corrupting the secondary. The result can be prevented by allowing resynchronization, or by overriding the safeguard.

User response: Allow the relationship to become

synchronized before deleting, or reissue the command with the -force flag to allow corruption of the secondary.

CMMVC7170E The Remote Copy relationship could not be created because the specified master VDisk is already a change VDisk for a different relationship.

Explanation: A master volume cannot be selected as a change volume of a Remote Copy relationship while currently defined for a different relationship.

User response: Choose a different master volume.

CMMVC7171E The Remote Copy relationship could not be created because the specified auxiliary volume is already a change volume for a different relationship.

Explanation: An auxiliary volume cannot be selected as a change volume of a Remote Copy relationship while currently defined for a different relationship.

User response: Choose a different auxiliary volume.

CMMVC7172E Enabling access to the secondary VDisk of the Remote Copy relationship could not be completed in a reasonable time.

Explanation: A timeout occurred before the task could be completed. The relationship is continuing to enable access, and will have a state of idling when access is enabled.

User response: Check the event log for any events to be resolved, and resubmit the task.

CMMVC7173E Enabling access to the secondary VDIs of the Remote Copy consistency group could not be completed in a reasonable time.

Explanation: A timeout occurred before the task could be completed. The consistency group is continuing to enable access, and will have a state of idling when access is enabled.

User response: Check the event log for any problems that need to be resolved, and resubmit the task.

CMMVC7174E The task cannot be completed because the other cluster is running a software version that is not recent enough.

Explanation: The software version of the one of the clusters is not supported.

User response: update the software version of the cluster.

CMMVC7175E Enabling access to the secondary VDisks of the Remote Copy consistency group could not be completed because the relationships in the group are not mutually consistent.

Explanation: The relationships in the consistency group must be mutually consistent before access to the secondary volumes can be enabled.

User response: Ensure that the relationships in the Remote Copy consistency group are mutually consistent.

CMMVC7176E The Remote Copy relationship could not be added to the consistency group because the cycling modes do not match.

Explanation: The cycling modes of the Remote Copy relationship and the consistency group to which it is being added must match.

User response: Ensure that the cycling modes match.

CMMVC7177E The Remote Copy relationship could not be added to the consistency group because the cycling periods do not match.

Explanation: The cycling periods of the Remote Copy relationship and the consistency group to which it is being added must match.

User response: Ensure that the cycling periods match.

CMMVC7178E The Remote Copy relationship could not be started in a reasonable time. It is now stopped.

Explanation: A timeout occurred before the task could be completed.

User response: Check the event log for any problems that need to be resolved, and resubmit the task.

CMMVC7179E The Remote Copy consistency group could not be started in a reasonable time. It is now stopped.

Explanation: A timeout occurred before the task could be completed.

User response: Check the event log for any problems that need to be resolved, and resubmit the task.

CMMVC7180E The Remote Copy relationship could not be started because no master change VDisk is defined.

Explanation: A master change volume must be defined for the Remote Copy relationship.

User response: Define a master change volume.

CMMVC7181E The Remote Copy relationship could not be started because no auxiliary change VDisk is defined.

Explanation: An auxiliary change volume must be defined for the Remote Copy relationship.

User response: Define an auxiliary change volume.

CMMVC7182E The Remote Copy consistency group could not be started because no master change VDisk is defined.

Explanation: A master change volume must be defined for the Remote Copy consistency group.

User response: Define a master change volume.

CMMVC7183E The Remote Copy consistency group could not be started because no auxiliary change VDisk is defined.

Explanation: An auxiliary change volume must be defined for the Remote Copy consistency group.

User response: Define an auxiliary change volume.

CMMVC7184E The task cannot be completed as the Remote Copy object is not stopped.

Explanation: The task cannot be completed as the Remote Copy object is not stopped.

User response: Stop the Remote Copy object.

CMMVC7185E The change VDisk could not be associated because the Metro Mirror or Global Mirror relationship has a VDisk on this cluster that is in a different I/O group.

Explanation: The I/O group of the change volume conflicts with an I/O group in the relationship with which an association was attempted.

User response: Ensure that no conflicting I/O groups exist.

CMMVC7186E The Remote Copy relationship was not created because the master VDisk is owned and has restricted use.

Explanation: The specified task cannot be performed while the master volume is in a file system or owned.

User response: Choose a different master volume, if the specified volume cannot be removed from the file system.

CMMVC7187E The Remote Copy relationship was not created because the auxiliary VDisk is owned and has restricted use.

Explanation: The specified task cannot be performed while the auxiliary volume is in a file system or owned.

User response: Choose a different auxiliary volume, if the specified volume cannot be removed from the file system.

CMMVC7188E The command failed because the master virtual disk (VDisk) is in a file system.

Explanation: The specified task cannot be performed on a master volume while it is in a file system.

User response: Choose a different master volume, if the specified volume cannot be removed from the file system.

CMMVC7189E The change VDisk could not be associated because it is in a file system.

Explanation: The specified change volume cannot be associated while it is in a file system.

User response: Choose a different change volume, if the specified volume cannot be removed from the file system.

CMMVC7203E The command failed because the hardware configuration of the local cluster is not compatible with the code of a partnered cluster.

Explanation: The hardware configuration of the local cluster is not compatible with the code of a partnered cluster. See the `chnodehw` explanation for more information.

User response: Make sure the hardware configuration and code level of all the clusters in a partnership are compatible before you create the partnership. Run `chnodehw` for diagnostic information.

CMMVC7205E The command failed because it is not supported.

Explanation: The command fails because it is not supported by the product.

User response: Review the documentation for the product to select an appropriate command.

CMMVC7206E The command failed because a parameter is not supported.

Explanation: The user entered a parameter which is not supported by the product they are using.

User response: Review the documentation and select

appropriate parameters for the product.

CMMVC7211E The command failed because it is not supported for image-mode MDisk.

Explanation: This error is returned by the `remove mdisk rmmdisk` command when issued against an MDisk that is backing an image-mode volume on a platform that does not support migration for image-mode volumes.

User response: If the image-mode volume is not required then use `rmvdisk` to delete the volume. This deletes the MDisk as well. If the user wants to migrate the image-mode volume's data to internal storage, do that using Volume Mirroring and then delete the image-mode volume copy.

CMMVC7218E The task cannot be initiated because an invalid license key was specified.

Explanation: The license key specified has not been recognized as a valid key.

User response: Check for typing errors and try again.

CMMVC7219E The task cannot be initiated because an invalid function id was specified.

Explanation: The function id specified in the command is invalid. There are a set number of features that can be activated. This message is shown when the user specifies an invalid function id.

User response: Use `lsfeature` to check which function id to specify.

CMMVC7220E The task cannot be initiated because an invalid function license key file path was specified.

Explanation: The file path specified for the license key file is invalid.

User response: Check for typing errors and try again.

CMMVC7221E The command has failed because the machine signature in the license does not match this machine.

Explanation: The license key provided is not valid for this storage system.

User response: Use a license key that has been generated using this enclosure's machine signature.

CMMVC7222E The command has failed because the function specified by the license key has not been recognized.

Explanation: The license key specified is for activating a function that is not supported on this level of firmware.

User response: update the storage system to a firmware level that supports the function and try again.

CMMVC7223E The command has failed because the trial specified has previously been used.

Explanation: Each function has a trial period that can only be redeemed once. The trial period on this machine has already been used.

User response: Purchase the full license in order to continue using this function.

CMMVC7224E The command has failed because the specified function is currently in use.

Explanation: A function can only be deactivated if it is no longer in use.

User response: Ensure that the function that you need to deactivate is no longer in use, and try again.

Reduce the number of FlashCopy targets to 64 or fewer before deactivating the function.

CMMVC7226E The command has failed because the license key file specified is not in a supported format

Explanation: The command has failed because the license key file specified is not in a supported format.

User response: Check that the correct file was uploaded and try again.

CMMVC7233W Easy Tier is active without a license for each enclosure.

Explanation: The system does not have enough licenses for Easy Tier.

User response: The user should correct their Easy Tier license entitlement.

CMMVC7234W FlashCopy is enabled without a license for each enclosure.

Explanation: The system does not have enough licenses for FlashCopy.

User response: The user should correct theirFlashCopy license entitlement.

CMMVC7235W Remote copy is enabled without a license for each enclosure.

Explanation: The system does not have enough licenses for remote copy.

User response: The user should correct their Remote copy license entitlement.

CMMVC7236W Multiple functions are enabled without a license for each enclosure.

Explanation: The system does not have enough licenses for multiple functions.

User response: The user should correct their license entitlement.

CMMVC7239E The node could not be added because the number of configured hosts exceeds the supported limit for the node type being added.

Explanation: One (or more) I/O groups has more hosts configured than is supported by the node hardware type being added.

User response: Either reduce the number of configured hosts or choose a different node hardware type to add.

CMMVC7240E The node could not be added because at least one host has more volumes mapped than is supported for the node type being added.

Explanation: One (or more) hosts has more volumes mapped than is supported by the node hardware type being added.

User response: Either reduce the number of volumes mapped or choose a different node hardware type to add.

CMMVC7241E System code update cannot start because a component firmware update is in progress.

Explanation: An attempt was made to initiate either a system code update or a firmware update while the system was updating the firmware of various hardware components. This update cannot be completed while a firmware download is in progress, so the request failed.

User response: The firmware download must complete before you can perform another update. Due to the dynamic nature of firmware downloads, you cannot follow the progress of the download. Wait approximately 10 minutes and retry the command. You might need to repeat this step several times. You can use the **svcinfo lsupdate** command to see whether the firmware download completed.

CMMVC7242E No help available for [%1].

Explanation: There is no help available for this command. [%1] shows the command for which help is not available.

User response: None.

CMMVC7243E The specified port mask cannot be applied because insufficient paths would exist for node communication.

Explanation: The localfcportmask port mask value that was specified would cause one or more nodes to lose contact with the system.

User response: Check zoning. Fix any port errors in the event log. Use the lsfabric CLI command to ensure that when the correct port mask is specified and applied, all nodes would still have two paths to contact every other node in the system.

CMMVC7248E You must use an 0x parameter for the -drivelba parameter.

Explanation: The parameter format for the `lsmdisklba` might be in error. See the description of the command to find an acceptable format.

User response: Try the command again using the parameter format described in the command documentation.

CMMVC7249E The name that you have entered is not valid. The name can contain letters, numbers, spaces, periods, dashes, and underscores. The name must not begin or end with a space. The name must not begin with a period.

Explanation: A period 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 period, does not begin or end with a space character, and does not contain any of the unsupported characters listed above, and resubmit the task.

CMMVC7300E The maximum number of volumes already exist.

Explanation: The maximum number of volumes have been created and one or more must be deleted.

User response: Destroy one or most of the volumes before more can be created.

CMMVC7301E The command failed because the volume would be smaller than the minimum size.

Explanation: A volume must be 1 MB or greater to be successfully created.

User response: Specify a capacity which is greater than the minimum supported capacity (1MB).

CMMVC7302E The command failed because insufficient extents are available.

Explanation: Too many volumes have been created. Contact IBM support for assistance.

User response: If possible, delete unused volumes and try again. If the error persists, the extent mapping must be defragmented.

Contact IBM support for assistance.

CMMVC7304E The command cannot be initiated because a spare is already configured.

Explanation: A configured spare already exists.

User response: Only new drives can be formatted. If a spare already exists, the command cannot be executed.

CMMVC7305E The command cannot be initiated because the drive fault is unrecoverable.

Explanation: A drive fault is preventing further actions.

User response: Replace the drive.

CMMVC7306E The command cannot be initiated because no array currently exists.

Explanation: No array has been created. No further action is possible.

User response: Create an array before using this command.

CMMVC7307E The command cannot be initiated because the battery slot number specified is invalid.

Explanation: Batteries occupy slot numbers 1 and 2 only for batteries.

User response: Select Slot 1 or 2.

CMMVC7308E The command cannot be initiated because an array already exists.

Explanation: An array exists that is prohibiting the desired action.

User response: Only a single array can be created. Either create new volumes on the same array, or remove all volumes and arrays, then recreate the array.

CMMVC7309E The command cannot be initiated because the specified raid level is unsupported on this platform.

Explanation: This platform supports either RAID 0 or RAID 5. No other RAID levels are supported.

User response: Specify either raid0 or raid5 as the raid level.

CMMVC7310E The command cannot be initiated because one or more drives is in the failed state.

Explanation: One or more drives has failed and is preventing most commands.

User response: Work through Directed Maintenance Procedures (DMP) associated with failed drives.

CMMVC7311E The command cannot be initiated because the number of drives is unsupported.

Explanation: The number of installed drives is incorrect for the desired configuration. RAID 0 requires 2, 4, 8 or 12 drives. RAID 5 requires 4, 8 or 12 drives.

User response: Remove or insert drives to obtain a supported configuration.

CMMVC7312E The command could not be initiated because one or more volumes are using the array.

Explanation: One or more volumes are already using the array. All volumes must be removed before removing the array.

User response: Remove all volumes before removing the array.

CMMVC7313E The specified array does not require recovery as it is not corrupt.

Explanation: The specified array is not corrupt and does not require repair.

User response: This command is only supported for corrupt arrays.

CMMVC7314E The specified array cannot be recovered due to failed drives.

Explanation: One or more drives in an array have failed and this is preventing recovery.

User response: Put the drives back in if they are missing.

CMMVC7315E The command cannot be initiated because the specified slot does not exist.

Explanation: The command was attempted on an invalid slot.

User response: Batteries have slots 1 and 2. Enclosures have slots 3-12.

Retry command with a slot that exists.

CMMVC7316E The specified array cannot be recovered due to failed drives.

Explanation: One or more failed drives are preventing the desired action. Too many errors.

User response: Change the volume SCSI Drive ID by enabling volume Open Access.

CMMVC7317E The command cannot be initiated because the specified PSU does not exist.

Explanation: A requested PSU does not exist.

User response: Select another PSU or install the requested PSU.

CMMVC7318E Open access is not enabled.

Explanation: Host mappings or logical host already exist.

User response: Delete host mappings or the logical host.

CMMVC7319E The air filter must be enabled to change the period.

Explanation: To change the timer for the air filter, the filter must currently be in use.

User response: Enable the filter timer, then try the command again.

CMMVC7320E Invalid port mask supplied.

Explanation: An invalid port mask was supplied.

User response: Review the command parameters and try again with corrected value.

CMMVC7322E Another volume uses this SCSI ID.

Explanation: The SCSI ID is already in use.

User response: Use a different SCSI ID or remove the other volume.

CMMVC7323E The command cannot be initiated due to a hardware fault.

Explanation: A hardware fault has occurred. Follow the Directed Maintenance Procedures (DMP).

User response: An unknown hardware fault exists. Follow the DMP to resolve the hardware fault.

CMMVC7324E The command cannot be initiated because the drive task is not supported.

Explanation: The drive does not support the command.

User response: None.

CMMVC7325E Open Access setting not changed because hosts exist.

Explanation: You cannot change the state of Open Access when hosts are defined.

User response: Remove the hosts and try the command again.

CMMVC7326E Cannot create logical host because open access is enabled.

Explanation: Open Access is preventing the creation of a logical host.

User response: Disable Open Access if host mapping access is desired

CMMVC7329E The maximum number of Infiniband GIDs for the host is already configured.

Explanation: The command has failed because the maximum number of Infiniband addresses has been exceeded.

User response: Remove unused hosts addresses and try again.

CMMVC7330E An invalid Infiniband GID was entered.

Explanation: The Infiniband address entered is not a valid address.

User response: Check your input and try again.

CMMVC7331E A supplied Infiniband GID is already assigned to another host.

Explanation: The Infiniband address configured is already assigned to a configured host.

User response: Check your input and try again.

CMMVC7332E The array does not exist.

Explanation: A flash memory array was not found by the Command Console LUN (CCL).

User response: Call IBM support for instructions to resume your failed upgrade.

CMMVC7334E The array could not be removed because it is initializing.

Explanation: A Remove Array operation was attempted while a newly created array was initializing.

User response: Allow the array to complete initialization before attempting to remove it.

CMMVC7336E The topology and speed are not compatible with the specified port.

Explanation: The user has entered a combination of speed and topology which is not compatible.

User response: Check your input and try again.

CMMVC7337E The command cannot be initiated because the encryption key was not found.

Explanation: The system could not locate a correct encryption key.

User response: Make sure to insert a USB drive that contains the correct encryption key into each node before you retry the command.

CMMVC7338E The command cannot be initiated because the encryption key is not valid.

Explanation: The encryption key that was provided could not be used.

User response: Make sure to insert a USB drive that contains the correct encryption key into each node before you retry the command.

CMMVC7339E The command cannot be initiated because the drive did not unlock.

Explanation: The system was unable to unlock the drive by using the encryption key that was provided.

User response: Make sure to insert a USB drive that contains the correct encryption key into each node before you retry the command.

CMMVC7340E The command cannot be initiated because the array is offline.

Explanation: Failure to create a volume because the array specified is offline.

User response: Bring the array online.

CMMVC7341E The update cannot be resumed at this time.

Explanation: The update cannot proceed because of hardware errors.

User response: Contact IBM support.

A service mode update may be necessary.

CMMVC7342E The array is already encrypted.

Explanation: The array is already encrypted.

User response: No response required.

CMMVC7343E A software upgrade is in progress and must finish before battery reconditioning can be started.

Explanation: Battery reconditioning cannot be performed while a software upgrade is in progress.

User response: Wait until the software upgrade has finished, and then restart battery reconditioning.

CMMVC7344E Battery reconditioning is not currently running.

Explanation: Battery reconditioning is not running, therefore it cannot be cancelled.

User response: No response required.

CMMVC7348E One of the batteries is currently reconditioning. Wait until reconditioning finishes.

Explanation: Battery reconditioning is already running on one of the batteries. Only one battery can be reconditioned at a time.

User response: Wait until the current battery reconditioning finishes before reconditioning another battery.

CMMVC7349E The other battery in the enclosure is not in a good and charged state.

Explanation: Both batteries in the enclosure must be in a good and charged state to maintain redundancy during the battery reconditioning process. The selected battery cannot be reconditioned because the other battery in the storage enclosure is not charged, not installed, or it has an unresolved error condition.

User response: Determine the condition of the other battery in the storage enclosure and take appropriate action. For example, install the battery if it is missing. If the battery has an error condition, check the event log and follow the recommended Directed Maintenance Procedure (DMP).

CMMVC7350E The selected battery cannot be reconditioned.

Explanation: The selected battery is not in a state in which it can be reconditioned. The battery is either not charged, not installed, or it has an unresolved error condition.

User response: Determine the condition of the selected battery and take appropriate action. For example, install the battery if it is missing. If the battery has an error condition, check the event log and follow the recommended Directed Maintenance Procedure (DMP).

CMMVC7351E Encryption cannot be disabled while an encrypted array exists.

Explanation: System encryption cannot be disabled if there is an encrypted array.

User response: Delete the encrypted array and then disable encryption.

CMMVC7352E The array cannot be encrypted because system encryption is disabled.

Explanation: The array cannot be encrypted because system encryption is disabled.

User response: Enable system encryption using the `chencryption` command and retry.

CMMVC7353E The command cannot be initiated because a rekey operation is in progress.

Explanation: The command cannot execute because there is a rekey operation in progress.

User response: Wait for the rekey operation to complete, and then re-enter the command. Use the `lsencryption` command to verify the status of the rekey operation.

CMMVC7354E Rekey operation failed.

Explanation: The rekey operation failed. Possibly to hardware errors or a missing USB drive.

User response: Please confirm that correct USB drive is inserted, and examine the event log to check for hardware errors. Call IBM Support if unable to determine the cause.

CMMVC7355E The command cannot be initiated because system encryption is not enabled.

Explanation: The command entered by the user requires system encryption to be enabled.

User response: Enable system encryption and then re-enter the command.

CMMVC7356E Unable to validate the key on the current USB drive.

Explanation: Validation failed. Insert the USB drive that contains the correct key for this system.

User response: Review the output of the `lseventlog` and `lsencryption` commands for additional information,

and take appropriate corrective action.

CMMVC7357E The command cannot be initiated because a drive has failed.

Explanation: The command cannot be initiated because a drive has failed.

User response: Resolve the problem that caused the drive to fail before retrying the operation. Review the event log for possible causes.

CMMVC7358E Rekey commit not allowed until keys have been copied.

Explanation: An attempt was made to commit a new key without making sufficient copies of the new key.

User response: Contact IBM Support.

CMMVC7359E Operation not applicable when encryption is disabled.

Explanation: Either key validation was attempted while system encryption was disabled, or an attempt was made to change an array from unencrypted to encrypted while system encryption was disabled.

User response: Enable system encryption and retry the command.

CMMVC7360E The command cannot be executed because the array is not online.

Explanation: The array has to be online to run recoverarray -validate.

User response: Use the lsarray command to check the value of the attribute parameter raid_status for the array. The command cannot be run if the array does not exist, or if the raid_status is **offline**, **degraded**, **syncing**, or **initiating**.

CMMVC7361E The command cannot be initiated because another operation is in progress on the array.

Explanation: Another operation on the array is already in progress.

User response: Wait until the current operation is completed before running another command.

CMMVC7362I The encryption copy tool should only be used with the assistance of support.

Explanation: The command should only be run with the assistance of IBM Support.

User response: Call IBM Support to continue.

Attention: Do not proceed without IBM Support assistance.

CMMVC7363E The system is unable to mount/unmount the USB drive.

Explanation: The command was unable to either mount or dismount the USB drive.

User response: Call IBM Support to continue.

Attention: Do not proceed without IBM Support assistance.

CMMVC7364E The system is unable to collect the information needed to provide the proper output from the lsencryption command.

Explanation: The command is unable to collect the output of the lsencryption command.

User response: Call IBM Support to continue.

Attention: Do not proceed without IBM Support assistance.

CMMVC7365E The system could not complete the key copy operation because it could not read the new key file.

Explanation: The new key file was not found on the USB drive.

User response: Insert the correct USB drive, or call IBM Support to continue.

CMMVC7366E The system could not complete the key copy operation because it could not read the current key file.

Explanation: The current key file was not found on the USB drive.

User response: If this is the first time a key copy operation has been run on a new system, then no response is required. Otherwise, ensure that the correct USB drive is inserted. Call IBM Support if you are unable to resolve the problem.

CMMVC7367E The key file is not valid.

Explanation: The key file is not valid.

User response: Ensure that the correct USB drive is inserted. Call IBM Support if unable to resolve the problem.

CMMVC7368I %1 additional copy has been made.

Explanation: Copy # %1 of the new key was made.

User response: No response required.

This is an informational message that states how many copies of the new key have been made.

CMMVC7369E No additional key copies were written to the USB drive; all required keys already exist on the drive.

Explanation: No additional keys were written to the USB drive because both the new and current keys already exist on the drive.

User response: Insert a new USB drive that does not contain any key files and retry the command, or call IBM Support to continue.

CMMVC7370W Only the new key file has been copied to the USB drive; the current key was already on the drive.

Explanation: Only the new key file has been written to the USB drive; the current key was already on the drive.

User response: The USB drive already contains the current key file. If the USB drive was expected to be empty, examine the files on the drive to determine the problem.

CMMVC7371E Only the current key file has been copied to the USB drive; the new key was already on the drive.

Explanation: Only the current key has been written to the USB drive; the new key was already on the drive. Call IBM Support to continue.

User response: Call IBM Support to continue. Do not proceed without IBM Support assistance.

CMMVC7372E The current key on the USB drive does not match the key that was previously read.

Explanation: The current key on the USB drive does not match the key that was previously read.

User response: Call IBM Support to continue. Do not proceed without IBM Support assistance.

CMMVC7373E The attempted write of the new key to the USB drive has failed.

Explanation: The attempted write of the new key to the USB drive has failed.

User response: Ensure that the correct USB drive is inserted. Call IBM Support if unable to resolve the problem.

CMMVC7374E The attempted write of the current key to the USB drive has failed.

Explanation: The attempted write of the current key to the USB drive failed.

User response: Ensure that the correct USB drive is

inserted. Call Support if unable to resolve the problem.

CMMVC7375E No keys were written to the USB drive. Do not proceed without Support assistance.

Explanation: No keys were written to the USB drive. Call Support to continue.

User response: Call Support to continue.

Attention: Do not proceed without Support assistance.

CMMVC7376E The command cannot be run because an encryption key has not been created.

Explanation: This command requires an encryption key, but the encryption state has not been properly set up.

User response: Use the encryption enablement procedure to create encryption keys, or contact IBM Support for assistance.

CMMVC7377E The command cannot be executed because the array does not exist.

Explanation: This command requires an encryption key, but the encryption state has not been properly set up.

User response: Create an array and then reattempt the command.

CMMVC7378E The command cannot be executed because the array is not encrypted.

Explanation: The array must be encrypted before running the command.

User response: Encrypt the array and then reattempt the command.

CMMVC7379E System encryption cannot be disabled because an encrypted array already exists.

Explanation: System encryption cannot be disabled while an encrypted array exists. To disable encryption, the encrypted array must be deleted, which will result in the loss of data.

User response: To disable encryption, delete the encrypted array, and then reattempt the command. Contact IBM Support for assistance before performing this step.

CMMVC7380E The command cannot be executed because an encrypted array exists.

Explanation: The command cannot be run because the array encryption cannot be disabled.

User response: No response is possible. The command is not allowed in this context.

CMMVC7381E The rekey failed because the current encryption key was not found on the USB drive.

Explanation: Explanation: The rekey operation could not be completed because the current encryption key was not found on the USB drive.

User response: Cancel the failed rekey operation, restore the current encryption key file to the USB drive, and then reattempt the operation.

CMMVC7382E The rekey failed because the new encryption key was not found on the USB drive.

Explanation: The rekey operation could not be completed because the new encryption key was not found on the USB drive.

User response: Cancel the failed rekey operation, restore the new encryption key file to the USB drive, and then reattempt the operation.

CMMVC7383E The rekey failed because the proposed new encryption key could not be generated.

Explanation: The rekey operation failed because the new encryption key could not be generated.

User response: Cancel the failed rekey operation and replace the USB drive before reattempting a new rekey operation.

If the command `chencryption -usb newkey -key prepare` was issued then the encryption key file may already exist on the USB drive.

Replace the USB drive with a blank USB drive and retry the command.

CMMVC7384E The rekey operation failed because one of the drives failed to prepare for the rekey.

Explanation: The rekey operation failed because the -key prepare option of the `chencryption` command failed on one of the drives.

User response: Review the event log and resolve the corresponding event.

CMMVC7385E The rekey operation failed because one of the drives failed to commit to the rekey.

Explanation: The rekey operation failed because the -key commit option to the `chencryption` command failed on one of the drives, causing the drive to fail to commit.

User response: Review the event log and resolve the corresponding event.

CMMVC7386E The rekey operation failed because one of the drives failed to cancel the rekey.

Explanation: The rekey operation failed because the -key cancel option to the `chencryption` command failed on one of the drives.

User response: Review the event log and resolve the corresponding event.

CMMVC7387E The command cannot be initiated because the rekey is not in an expected state.

Explanation: The command cannot be initiated because the rekey is not in an expected state.

User response: Issue the appropriate command based on the current rekey state.

A commit operation can only be performed when the rekey state is `prepare` or `commit_failed`.

A cancel operation can only be performed when the rekey state is `prepare_complete`, `prepare_failed` or `cancel_failed`.

CMMVC7388E Rebuild options are not supported for this array RAID level.

Explanation: The command cannot be initiated because the rekey is not in an expected state.

User response: Issue the command against a RAID 5 array. Do not target a RAID 0 array.

CMMVC7395E The command cannot be executed because the enclosure is not a part of the cluster.

Explanation: You tried to create an array with a flash enclosure that is not managed.

User response: Change the flash enclosure to a managed type by using the following command:
`chenclosure -managed yes enclosureId`

where *enclosureId* is the ID of the enclosure as displayed in the `lsenclosure` command. Then, retry the command.

CMMVC7403E The command cannot be initiated because it is not supported for the specified enclosure type

Explanation: The specified enclosure was not the correct type for this command.

User response: Review the command documentation or specify an alternative enclosure.

CMMVC7404E The command cannot be initiated because it is not supported for the specified drive type

Explanation: The specified drive was not the correct type for this command.

User response: Review the command documentation or specify an alternative drive.

CMMVC7405E The command cannot be initiated because it is not supported for the specified array type

Explanation: The specified array was not the correct type for this command.

User response: Review the command documentation or specify an alternative array.

CMMVC7406E The command failed because the requested volume size is too large. Reduce the array's reserved capacity to free space.

Explanation: The `mkvdisk -size` and `chvdisk -size` commands cannot encroach on the array space that is reserved for performance enhancement.

User response:

- Specify a smaller volume size.
- Reduce the amount of space in the array that is reserved to increase performance.

CMMVC7407E The array has insufficient free space to reserve for performance.

Explanation: The array does not have sufficient free space reserved for performance.

User response: Complete one of the following tasks:

- Use the `mkarray -reservesize` command to reserve a size that is smaller than that of the array.
- Delete volumes to increase available space in the array.

CMMVC7408E The command has failed because one or more of the parameters entered is not valid for the specified enclosure type.

Explanation: Some parameters are only valid when used with certain types of enclosure.

User response: Review the command syntax for the type of enclosure specified or specify an alternative enclosure.

CMMVC7409E The command cannot be initiated because the port is not online.

Explanation: `chportip` must be used on a port that is online.

User response: Use `lsportip` to see the state of the port. If the state is listed `asoffline`, the port is offline.

CMMVC7410E The command has failed because one or more of the parameters entered is not valid for the specified drive type.

Explanation: Some parameters are only valid when used with certain types of drives.

User response: Review the command syntax for the type of drive specified or specify an alternative drive.

CMMVC7411E The command has failed because one or more of the parameters entered is not valid for the specified array type.

Explanation: Some parameters are only valid when used with a certain type of array.

User response: Review the command syntax for the type of array specified or specify an alternative array.

CMMVC8000E Cannot execute on an active node.

Explanation: This operation cannot be executed on an active node.

User response: Choose an inactive node, or deactivate the node before trying the operation again.

CMMVC8001E Cannot execute on a candidate node.

Explanation: This operation cannot be executed on a candidate node.

User response: Choose another node and try the operation again.

CMMVC8002E Cannot execute on a service-state node.

Explanation: This operation cannot be executed on a service-state node.

User response: Choose another node and try the operation again.

CMMVC8003E Cannot execute on a node in a cluster-recovery state.

Explanation: This operation cannot be executed on a node in a cluster-recovery state.

User response: Choose another node and try the operation again.

CMMVC8004E Cannot execute on a node with a location error.

Explanation: You cannot process this operation on a node with a location error.

User response: Fix the error or choose another node and try the operation again.

CMMVC8005E Cannot execute on a node displaying hardware errors.

Explanation: You cannot execute this operation on a node displaying hardware errors.

User response: Fix the errors or choose another node and try the operation again.

CMMVC8006E Cannot execute on a node displaying errors.

Explanation: You cannot execute this operation on a node displaying errors.

User response: Fix the errors or choose another node and try the operation again.

CMMVC8007E Cannot execute on a node which is charging.

Explanation: You cannot execute this operation on a node that is charging its batteries.

User response: Wait until the batteries have finished charging and try the operation again.

CMMVC8008E Cannot execute on this SAN Volume Controller node.

Explanation: You cannot execute this operation on this SAN Volume Controller node.

User response: Choose another node and try the operation again.

CMMVC8009E Cannot execute on a node canister.

Explanation: You cannot execute this operation on a node canister.

User response: Choose an appropriate target and try the operation again.

CMMVC8010E Not from a USB stick.

Explanation: You cannot execute this operation from a USB stick.

User response: Change location to an appropriate place and try the operation again.

CMMVC8011E Version too high for this client.

Explanation: You cannot complete this operation during a manual update.

User response: Wait until the manual update finishes and try the operation again.

CMMVC8012E The operation did not complete in the time allowed.

Explanation: The operation did not complete in the time allowed.

User response: Set the time allowed for operations to complete appropriately. Determine if another operation or error caused the problem.

CMMVC8013E Incompatible parameters set.

Explanation: Parameters provided are mutually exclusive.

User response: Set appropriate parameters and try the operation again.

CMMVC8014E Bad parameter value.

Explanation: The command was entered with an unparseable ip address, WWNN, or some other unknown parameter value.

User response: Set appropriate parameters and try the operation again.

CMMVC8015E Failed.

Explanation: The enclosure midplane has a cluster ID set or the next cluster ID on the midplane is corrupt or invalid.

User response: Fix the problem and try the operation again.

CMMVC8016E Node would be clustered if not in service state.

Explanation: Cannot run because the node will be in cluster when it exits service.

User response: Fix the problem and try the operation again.

CMMVC8017E Info value not recognised.

Explanation: Info value not recognized.

User response: Use a valid info value and try the operation again.

CMMVC8018E Provided buffer was too small.

Explanation: Provided buffer was too small.

User response: Increase the buffer size.

CMMVC8019E Task could interrupt IO and force flag not set.

Explanation: Running on an active node could impact I/O.

User response: Wait until the node is inactive before performing this task.

CMMVC8020E Attempting to create cluster while there is a stored cluster ID.

Explanation: An attempt was made to create a cluster while the control enclosure or the node has a stored cluster ID.

User response: Change the cluster ID or chose a different control enclosure or node.

CMMVC8021E Invalid panel name.

Explanation: Invalid panel name given in the parameter.

User response: Use a valid panel name.

CMMVC8022E New cluster created but node cannot leave service state.

Explanation: New cluster created but node cannot leave service state. The battery might be charging or some other service task running.

User response: Wait until all service tasks complete.

CMMVC8023E Partner node is clustered.

Explanation: Cannot execute because the partner node is clustered.

User response: Take the partner node out of the cluster, or choose a different node.

CMMVC8024E Gateway or subnet/prefix required.

Explanation: A Gateway or subnet/prefix required.

User response: Use a Gateway or subnet/prefix.

CMMVC8025E DHCP Failed.

Explanation: DHCP Failed.

User response: Try operation again.

CMMVC8026E No suitable donor.

Explanation: No suitable donor.

User response: [User response needed].

CMMVC8027E T3 prepare failed.

Explanation: T3 prepare failed.

User response: [User response needed].

CMMVC8028E T3 prepare incomplete.

Explanation: T3 prepare not complete.

User response: [User response needed].

CMMVC8029E T3 execute failed.

Explanation: T3 execute failed.

User response: [User response needed].

CMMVC8030E Another instance of this command is already running.

Explanation: Cannot execute because another instance of this command is already running.

User response: Wait for the other instance of the command to complete.

CMMVC8031E file not found.

Explanation: Required/provided file is not on the file system in the expected location.

User response: Locate the missing file.

CMMVC8032E The specific update package cannot be installed over the current version.

Explanation: Cannot install this code over current version; cannot update to this version; code is already at this level.

User response: Make sure you are installing the right version.

CMMVC8033E Password reset disabled.

Explanation: The password reset function is disabled.

User response: Enable the password reset function or contact your system administrator.

CMMVC8034E A compulsory parameter is missing.

Explanation: Cannot execute because required parameter has not been provided.

User response: Provide the required parameter.

CMMVC8035E The service assistant CLI is not ready - try again.

Explanation: The Service CLI interface is not ready/running yet.

User response: Wait a few moments and try again.

CMMVC8036E No help available.

Explanation: There is no available help.

User response: Contact IBM Support.

CMMVC8037E A required parameter is missing.

Explanation: A required parameter is missing.

User response: Rerun the operation with the required parameter set.

CMMVC8038E Required parameters are missing.

Explanation: Required parameters are missing.

User response: Rerun the operation with the required parameters set.

CMMVC8039E The [%1] parameter is missing its associated arguments.

Explanation: The parameter is missing an argument.

User response: Rerun the operation with the required arguments for the parameter.

CMMVC8040E [%1] is not a supported parameter.

Explanation: The parameter is not supported.

User response: Rerun the operation using a supported parameter.

CMMVC8041E [%1] is not a valid command line option.

Explanation: The command given does not exist.

User response: Use an existing command.

CMMVC8042E Invalid or inconsistent arguments.

Explanation: Invalid or inconsistent arguments. For example, the argument at the end is not a recognized panel id.

User response: Use valid and consistent arguments.

CMMVC8043E This command can only be run by the superuser.

Explanation: Cannot execute because user is not superuser.

User response: Have a superuser run the command.

CMMVC8044E Command completed successfully.

Explanation: Command completed successfully. This message is only used within lscmdstatus.

User response: None.

CMMVC8045E Authentication failed.

Explanation: Authentication failed.

User response: [Need user response].

CMMVC8046E Partner node has lost cluster data.

Explanation: Partner node has lost cluster data.

User response: [Need user response].

CMMVC8047E Not a valid ssh key.

Explanation: Argument is not a valid ssh key.

User response: Use a valid ssh key for the argument.

CMMVC8048E Invalid file permissions.

Explanation: Cannot execute the argument. The file has invalid file permissions.

User response: Set valid file permissions.

CMMVC8049E Invalid cluster name.

Explanation: User provided an invalid cluster name.

User response: Use a valid cluster name.

CMMVC8050E Files cannot be unpacked from the update package. The system must unpack them.

Explanation: Possible causes are:

1. A Duff boot drive or sector.
2. Full /upgrade, /tmp, or /upgrade.
3. An invalid package.

User response:

1. Clear all dumps and retry the installation.
 2. Reboot the node and retry the installation.
-

CMMVC8051E Utilities package installed.

Explanation: The utilities package installed successfully.

User response: None.

CMMVC8052E The signature of the utility package has failed verification.

Explanation: This can be caused by the following issues:

1. The package has become corrupted.
2. The package is not a valid IBM utility.
3. The system clock on the node is long out-of-date and the package signature is too far into the future.

User response:

1. Ensure that the installation package is complete and retry the installation.
 2. Verify that the utility was provided by IBM support personnel.
 3. Use chsystemtime Change the system clock to accurately reflect the date.
-

CMMVC8053E The specific update package cannot be installed on this hardware.

Explanation: The software is incompatible with the hardware level.

User response: Ensure that you have downloaded the appropriate package for the hardware you are updating.

CMMVC8054E The update failed as the package is missing files.

Explanation: The package can be missing files due to a bad upload.

User response: Validate that the package was downloaded or uploaded properly and try to run the update again.

CMMVC8055E The command cannot be run because the node is busy.

Explanation: The node can only run one task command at a time, or the enclosure firmware is being updated so the command cannot be run at this time.

User response: Wait for the task to complete and run the command again.

CMMVC8056E An error occurred in communicating with the target node.

Explanation: The error can be caused by:

1. A fault in one partner node is preventing another partner node from seeing it.

2. The Fibre Channel network is congested or faulty, and packages are failing to transfer.

User response:

1. If you are using SAN connectivity, check your fabric to ensure that all the nodes in a cluster have clear paths to one another.
 2. Ensure that the target node is online and does not have a hardware or location error (if it is Enclosure based).
 3. Ensure that both the source and target nodes see each other in lsservicenodes. If they do not see each other, a path is missing.
 4. If trying the first three steps does not correct the problem, work directly on the node in question rather than from the partner node. Log onto the target node's service assistant and run the operation locally.
-

CMMVC8057E Files cannot be copied to a node which is not part of the source cluster.

Explanation: Files cannot be copied to a node which is not part of the source cluster using cpfiles.

User response: Choose a node that is part of the source cluster.

CMMVC8058E Cannot create file, file already exists.

Explanation: The ssh key file you are trying to create already exists.

User response: None.

CMMVC8059E The update package supplied cannot be installed using service state whilst maintaining the cluster configuration on this node. To maintain the cluster configuration on this node, this package can only be installed with applysoftware or pacedccu mode. This package can be installed in service state using the -ignore flag, however the cluster state will be destroyed and cluster configuration will be lost from the node.

Explanation: A software update to this level cannot be issued when it uses service state without using the **-ignore** flag. The **-ignore** flag removes the cluster configuration from the node. If you want to maintain the cluster configuration, you must install the package by using the automated **applysoftware** command or by manual update.

Attention: Be careful when you use the **-ignore** flag. You might see adverse consequences to the data that is processed.

User response: Use the correct procedure to update the code package.

CMMVC8060E DHCP fallback is not supported on this platform.

Explanation: You tried to set the service IP via DHCP with fallback enabled. This platform does not support the fallback option.

User response: Set the service IP via DHCP without fallback enabled. This product does not support the fallback option.

CMMVC8061E The enclosure does not support setting a machine part number.

Explanation: Cannot execute because there is no -machinepartnum on this system.

User response: None.

CMMVC8062E The machine type and model is not valid for this enclosure.

Explanation: The machine type and model is not valid for this enclosure.

User response: Use a valid machine type and model with the command.

CMMVC8063E The machine part number is not valid for this enclosure.

Explanation: The machine part number is not valid for this enclosure.

User response: Use a valid machine part number with the command.

CMMVC8064E The machine part number and machine type and model are not valid for this enclosure.

Explanation: The machine part number and machine type and model are not valid for this enclosure.

User response: Use a valid machine part number and machine type and model with the command.

CMMVC8065E An enclosure property has already been set and cannot be modified.

Explanation: An enclosure property has already been set and cannot be modified. A valid VPD exists and will not be overwritten.

User response: If the machine serial number is not 00000000 or is not the same as the serial number stored on one of the boot drives, then the system board must be replaced again with a clean FRU that has a machine serial number of 00000000 stored on it. .

CMMVC8066E A new enclosure VPD field does not match the node copy.

Explanation: A new enclosure VPD field does not match the node copy. An unexpected value was given.

User response: Rerun the command with the correct value in the enclosure VPD field.

CMMVC8070E The specified IP addresses are not in the same sub-network.

Explanation: The IP address parameter values in satask.txt and cfgtask.txt that must be in the same sub-network are not all in the same sub-network. i.e. The bit wise AND of an IP address with the subnet mask is not the same as the bit wise AND of another IP address with the subnet mask.

User response: Check and correct the parameters in satask.txt and cfgtask.txt or use InitTool to create new valid satask.txt and cfgtask.txt files on the USB flash drive to start the initial setup of the system.

CMMVC8071E The specified IP address is already in use.

Explanation: An arping of the IP address has received a reply from an IP host on the network that is already using this IP address from satask.txt.

User response: If you cannot stop the other IP host from using that IP address then allocate another IP address and use InitTool to create new valid satask.txt and cfgtask.txt files on the USB flash drive to start the initial setup of the system.

CMMVC8072E The above parameters are invalid or missing from cfgtask.txt.

Explanation: The required parameters shown above this error message are invalid or missing from cfgtask.txt.

User response: Check and correct the parameters in cfgtask.txt or use InitTool to create new valid satask.txt and cfgtask.txt files on the USB flash drive to start the initial setup of the system.

CMMVC8085E The command failed because the node does not support technician port functionality.

Explanation: Some older models do not provide a technician port.

User response: To connect a workstation directly to a node without a technician port, connect the workstation using Ethernet port 1 or 2 and an IP address compatible with the system IP or service IP of the node. If the IP configuration of these ports is not known, use a USB flash drive inserted into the USB port of the node.

CMMVC8086E The command failed because the node already has a dedicated technician port.

Explanation: When a node has a dedicated technician port, it is not possible to configure another port as a technician port. The dedicated technician port is permanently enabled.

User response: If you need to connect a workstation directly to the node for maintenance, use the dedicated technician port.

CMMVC8087E The command failed because the file specified is a valid USVNID file but it is for a different node.

Explanation: An activation file was provided that is in the correct format but has an incorrect unique ID (obtained from the node during the installation process).

User response: Verify that the correct file was downloaded from the IBM web server. You might need to repeat the steps that you took to generate this file, making sure that you use the correct node ID.

CMMVC8088E The command failed because the file specified is not a valid USVNID file.

Explanation: An activation file was provided that is in the wrong format.

User response: Verify that the correct file was downloaded from the IBM web server. You might need to repeat the steps that you took to generate this file.

CMMVC8261E The command failed because the hardware configuration of the local cluster is not compatible with the software of a partnered cluster.

Explanation: The software version of the local cluster is newer than the software version of a partnered cluster, and additional hardware has been enabled that is not supported by the older software on the partnered cluster.

User response: Either update the software on the partnered cluster, turn off the new hardware on the local cluster, or stop the remote copy relationship with the remote cluster. Use the CLI command `chnodehardware -legacy` to disable hardware that is not supported by older software versions.

CMMVC8262E Cannot remove the latest I/O group from the volume access set.

Explanation: Running this command removes access to all the I/O groups in the access set.

User response: Before you run the command again, either modify the I/O group list so that it does not

include all the I/O groups that provide access to the volume or add more I/O groups to the access set.

CMMVC8263E The command failed because the volume is associated with a file system and your requested action can not be completed under your current user role.

Explanation: You are attempting to complete an action on a volume that is associated with a file system. However, you do not possess the required role for file system actions.

User response: Execute command through volume command.

CMMVC8264E The command failed because the volume associated with a file system and only the real capacity of a compressed file system volume can be changed.

Explanation: You are attempting to resize a volume that is associated with a file system. However, you can only resize the real capacity of a file system volume if it is compressed.

User response: The command cannot be completed on this volume. It will only succeed with a volume that is not associated with a file system or a with a compressed file system volume.

CMMVC8265E The action failed because the specified port is for management only.

Explanation: The action failed because the specified port is for management only.

User response: Try another port which is not marked as `management_only` in the output of `lspportip`.

CMMVC8266E The action failed because the specified port is not installed.

Explanation: The action failed because the specified port is not installed.

User response: Use a port that is shown in the output of `lspportip` but which is not marked as `management_only`.

CMMVC8267E The attempt to add the node to the system failed because the node is being manually updated. The system must be prepared for update by using the same package that is currently installed on the node that is being updated.

Explanation: The manual update requires that you prepare the system for the update first. Manual update mode (the `-pacedccu` parameter) is in use and an attempt is being made to add a node to the system at a

newer code level. However, one of the following situations occurred:

- The code level of the node that is being added requires that you prepare the system for update before you attempt to use the **addnode** command
- The system was prepared at a different code level than the version of software on the node that is being added

User response: Prepare the system for update with the same software version as the manual update node that you are adding. To do so, complete one of the following actions:

- Cancel the update and prepare the system with the correct package.
- Install a version of the code on the new node that matches the system version.

CMMVC8268E The attempt to prepare the cluster for update has failed because the cluster has already been prepared with a different package level. The update must first be cancelled before reattempting the update.

Explanation: The user has prepared the cluster for update with one level package and then attempted to prepare the update with a different, higher level package or attempted to automate the update to a higher level after the cluster was prepared with a lower level package.

User response: The user must cancel the current update that is in progress and re-prepare with the desired update package.

CMMVC8269E The attempt to prepare the cluster for update has failed because the previous update is in prepare_failed state. The previous update must first be aborted before reattempting the update.

Explanation: The current status of `lsupdate` reports the update as `prepare_failed`. This is an indication the user has already attempted to prepare an update, or started an update and in either scenario, the prepare has failed due to offline volumes. The cache flush has failed.

User response: The user needs to correct the error that caused the prepare to fail. Offline volumes are the most likely cause, also node resets might cause a prepare failure. Stop updates with prepare failures with the `applysoftware -abort` command, then reattempt the update.

CMMVC8270E The `applysoftware` prepare timed out because an attempt to make the volume cache empty took too long. The command will be completed asynchronously. Use `lsupdate` to monitor the progress.

Explanation: The `applysoftware` prepare timed out because an attempt to make the volume cache empty took too long. The command will be completed asynchronously. Use `lsupdate` to monitor the progress. The state will be reported as prepared when successfully completed.

User response: Wait until the prepare completes and `lsupdate` reports prepared.

CMMVC8272E Access iogrp parameter not valid when creating a file system volume.

Explanation: Access iogrp parameter not valid when creating a file system volume.

User response: Rerun `mkvdisk` without the `-accessiogrp` parameter or without the `-filesystem` parameter.

CMMVC8273E Host cannot be removed as there is a volume that is accessible from multiple iogrps including one of the iogrps specified.

Explanation: If a volume is mapped to a host, it must be mapped in all of the iogrps in which it is accessible. The `rmhostiogrp` command fails if it would leave the system in this state.

User response: Use `lshostvdiskmap` to find the list of volumes that are mapped to the host in multiple iogrps. Then for each one either a) remove the host/volume mapping or b) remove the iogrp that the host is being removed from the volume's access iogrp set.

CMMVC8274E The entry in the event log cannot be fixed because the given sequence number is out of range.

Explanation: The event log entry sequence number must be in the range 100 to 9,999,999 inclusive.

User response: Supply a valid event log entry sequence number in the range 100 to 9,999,999 inclusive.

CMMVC8275E An entry with the given sequence number cannot be found in the event log.

Explanation: The fix request has failed because an entry with the given sequence number cannot be found in the event log.

User response: Supply a sequence number of an entry that exists in the event log.

CMMVC8276E The entry in the event log cannot be fixed because it has expired or is in the monitoring state.

Explanation: The entry in the event log cannot be fixed because it has expired or is in the monitoring state.

User response: Expired and monitoring entries in the event log cannot be fixed.

CMMVC8277E The MTM format must be XXXX-YYY where X is a numeric value, and Y is numeric or upper case character.

Explanation: The user has attempted to change the MTM but provided an incorrect format.

User response: Reissue the command with the MTM with the correct format. The format must be XXXX-YYY where XXXX are numeric values and YYY are alphanumeric characters. Any alphabetic characters must be upper case.

CMMVC8278E The volume is accessible through more than one I/O group, and the host being mapped to the volume does not support volumes being mapped from multiple I/O groups.

Explanation: The volume is accessible through more than one I/O group, and the host being mapped to the volume does not support volumes being mapped from multiple I/O groups.

User response: Choose a different host or volume to map.

CMMVC8279E The volume is accessible through more than one I/O group, and the host being mapped to the volume has an iSCSI name. iSCSI hosts do not support volumes being mapped from multiple I/O groups.

Explanation: The volume is accessible through more than one I/O group, and the host being mapped to the volume has an iSCSI name. iSCSI hosts do not support volumes being mapped from multiple I/O groups.

User response: Choose a different host or volume to map.

CMMVC8280E The host has at least one volume mapped which is accessible through more than one I/O group, and the port being added is from a host system which does not support volumes being mapped from multiple I/O groups.

Explanation: The host has at least one volume mapped which is accessible through more than one I/O group, and the port being added is from a host system

which does not support volumes being mapped from multiple I/O groups.

Note: This error does not apply to release 7.6.0 and up.

User response: Choose a different port to add to the host.

CMMVC8281E The host has at least one volume mapped which is accessible through more than one I/O group, and the port being added is from a host with an iscsi name. Iscsi hosts do not support volumes being mapped from multiple I/O groups.

Explanation: The host has at least one volume mapped which is accessible through more than one I/O group, and the port being added is from a host with an iscsi name. Iscsi hosts do not support volumes being mapped from multiple I/O groups.

Note: This error does not apply to release 7.6.0 and up.

User response: Choose a different port to add to the host.

CMMVC8282E At least one host mapped to the volume does not support volumes being mapped from multiple I/O groups.

Explanation: At least one host mapped to the volume does not support volumes being mapped from multiple I/O groups.

User response: Unmap the host which does not support access from multiple I/O groups.

CMMVC8283E At least one host mapped to the volume has an iscsi name. Iscsi hosts do not support volumes being mapped from multiple I/O groups.

Explanation: At least one host mapped to the volume has an iscsi name. Iscsi hosts do not support volumes being mapped from multiple I/O groups.

User response: Unmap the host which does not support access from multiple I/O groups.

CMMVC8290E The action failed because the requested combination of notification settings is not permitted.

Explanation: The valid combinations of notification settings are info+warning+error and warning+error.

User response: Reissue the command using a valid combination of notification settings.

CMMVC8292E The command failed because a parameter is not supported.

Explanation: The user typed a command that is supported, but used a parameter that is not supported on this platform -- the parameter is supported on other platforms. For example the -mdisk parameter on the mkmdiskgrp command is not supported on this platform, but is supported on other platforms.

User response: Check the syntax in the CLI guide; check the system is the intended one.

CMMVC8293E The command failed because it is not supported for image-mode volumes.

Explanation: This error is returned by a migration command when issued against an image-mode volume on a platform that does not support migration for image-mode volumes.

User response: Perform these steps:

1. Use Volume Mirroring to change the volume's storage pool.
 2. Add a volume copy in the desired storage pool.
 3. Sync the volume.
 4. Delete the first copy.
-

CMMVC8294E The command failed because the FlashCopy feature is not active and the maximum number of FlashCopy target volumes already exist.

Explanation: The user has already configured the maximum number of FlashCopy targets that are allowed without activating the FlashCopy licensed feature.

User response: Delete some FlashCopy targets or activate the FlashCopy licensed feature. Try the command again.

CMMVC8295E The command failed because a licensed feature is not activated.

Explanation: This command requires a licensed feature to be activated before it is used.

User response: Activate the licensed feature and try again.

CMMVC8297E The drive cannot be managed because it has become unreachable.

Explanation: This message occurs if the user tries to change the use of the drive too soon after the drive has been inserted or the enclosure has been connected for the first time. It can also occur because of a hardware fault or if the user tries to change the use of a drive that has been removed from the system but is already managed.

User response: Wait and try again. If this doesn't work after 10 minutes, replace the drive. A drive that has been removed can still have its use changed to unused.

CMMVC8299E The system ran out of temporary resource while opening the file.

Explanation: The temporary directory used to unpack the file cannot be opened.

User response: In the unlikely event of this error occurring, schedule either a node reboot or a node rescue maintenance task.

Once the maintenance task completes and the node is online, try the command again.

CMMVC8309E The task cannot be initiated because some of the specified drives have an unsupported drive technology.

Explanation: The drive technology is the value of the `tech_type` field that is returned by the `svcinfo lsdrive` command. The following values are supported:

- tier0_flash
- tier1_flash
- tier_enterprise
- tier_nearline

Specifying a drive with any other technology type results in this error.

User response: Use the `lsdrive` command to determine which drives have an unsupported drive technology.

Repeat the command but do not include the drive ID of any drive with an unsupported drive type.

CMMVC8311E The system cannot open the file.

Explanation: The file specified after the -file option cannot be opened.

User response: Refer to the documentation for the update file to ensure that the update file is correct.

Obtain a new copy of the correct package file, copy it to the system, then run the command again.

CMMVC8313E The specified file cannot be used because it contains too many drive software images.

Explanation: The drive package contains more files than can be unpacked.

User response: Use a new package file and try the command again.

CMMVC8314E The system cannot read drive software from the specified file because the file contains an invalid drive firmware version string.

Explanation: The drive package has probably become corrupt or has been made incorrectly.

User response: Refer to the documentation for the update file to ensure that the update file is correct.

Obtain a new copy of the correct package file, copy it to the system, then run the command again.

CMMVC8315E The system cannot read drive software from the specified file.

Explanation: An internal error has occurred extracting the drive firmware from the drive package.

User response: Refer to the documentation for the update file to ensure that the update file is correct.

Obtain a new copy of the correct package file, copy it to the system, then run the command again.

CMMVC8316E The system can only program the FPGA of one drive at a time.

Explanation: More than one drive has been specified and the -type option has been set to fpga.

User response: Change the command to only specify a single drive ID because svctask applydrivesoftware -type fpga only supports one drive by one command.

CMMVC8317E The system cannot apply the task because an earlier drive update task is still ongoing.

Explanation: There is an existing svctask applydrivesoftware already running, and only one is allowed at a time.

User response: Try the command again when there is no drive update task in progress. Either use the lsdriveupgrade progress command to determine when the updates will be completed, or use svctask applydrivesoftware -cancel to cancel the current drive update task.

CMMVC8318E The system cannot apply drive software to the specified drives because some of them are not online.

Explanation: An svctask applydrivesoftware command has been issued, but some drives that have been specified are not online.

User response: Carry out problem determination to bring the drives online, then repeat the command. Alternatively, repeat the command but do not specify the drive ID of any drive that is not online.

CMMVC8319E Some of the specified drives are offline. The system cannot apply drive software to an offline drive, even if the -force option is specified.

Explanation: The svctask applydrivesoftware command was issued, but some specified drives are offline.

User response: Ensure that the specified drives are in either online or degraded state, then reissue the command.

CMMVC8320E None of the drives can be scheduled for drive software upgrade.

Explanation: None of the specified drives are in a suitable state to apply drive firmware updates.

User response: Ensure that the drive states meet all the requirements to permit drive software download before repeating the command.

CMMVC8321E The '-all' or '-drive' option is required to specify the drive(s) you want to upgrade.

Explanation: An attempt was made to apply software to one or more drives without specifying a drive.

User response: Retry the applydrivesoftware command and this time specify either the -all option or the -drive option followed by one or more colon-separated drive IDs.

CMMVC8323E The task cannot be applied to unused drives when multiple drives are specified.

Explanation: Some of the specified drives cannot be upgraded because the drives are currently unused by the system.

User response: Repeat the command, but do not include the drive ID of any drive that is currently unused.

CMMVC8324E There are no drive software upgrades scheduled.

Explanation: The command is not ongoing while user input -cancel option.

User response: No action is required.

CMMVC8325E None of the specified drives needed to be upgraded or downgraded.

Explanation: Every requested drive firmware is up-to-date in default. If the package is old, or contains no newer images for the drives on the package, the command cannot pass the level check.

User response: Ensure the package is correct or use -allowreinstall or -allowdowngrade option, then repeat the command.

CMMVC8326E The task cannot be initiated because the drive use changed.

Explanation: The command is stopped if user change the use of the drive, because some drive may be changed to "unused" while the command is ongoing.

User response: Check the use of the drive that you specified on the command line. If it is still appropriate to upload new firmware to the drive, repeat the command.

CMMVC8327E The -allowreinstall and -allowdowngrade options cannot be used with option -type fpga.

Explanation: When we applydrivesoftware fpga type drives, we don't permit to reinstall or downgrade drives.

User response: To download drive FGPA software, repeat the command but ensure the -allowreinstall and -allowdowngrade options are omitted.

CMMVC8328E No package file is specified or invalid package file name is used.

Explanation: User input invalid package file name.

User response: Input right package file name, and repeat the command.

CMMVC8329E The task cannot be initiated because downloading to one or more drives could cause volumes to go offline. Force is required.

Explanation: With any drive software update there is a risk that the drive might become unusable. If the drive is a member of a RAID0 array, consider whether to introduce additional redundancy to the protect the data on that drive.

User response: If the drive is not a member of a RAID0 array, fix any errors in the event log that relate to the array. When the drive is a member of an array with sufficient redundancy, repeat the command. Alternatively, consider using the '-force' option.

CMMVC8330E The %1 of %2 %3 in backup is %4; should be %5

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.

CMMVC8331E The command failed because at least one of the specified MDisks can only be used in image mode.

Explanation: Image mode only MDisk is forbidden to be added into the storage pool.

User response: Do not include MDisks that can only be used in image mode in a storage pool.

CMMVC8332E No MDisks were removed from the storage pool because at least one of the specified MDisks can only be used in image mode.

Explanation: Cannot remove image mode only MDisk with this command.

User response: Use rmvdisk or rmvdiskcopy to remove the MDisks which could only be used in image mode.

CMMVC8333E The task failed because it is not supported for image mode only MDisks.

Explanation: Cannot migrate data away or to image mode only MDisk with this command.

User response: Use Volume Mirroring to migrate data from or to an image mode only MDisk.

CMMVC8336E Site was not specified. Site must be specified because topology is stretched.

Explanation: Site was not specified. Site must be specified because topology is stretched.

User response: Identify the site of the new node and resubmit the command with -site flag.

Or, change system topology.

Note: Changing the topology will disable the DR feature.

CMMVC8337E Site specified not valid. System topology is stretched and the other member has been configured to the same site.

Explanation: Site specified not valid. System topology is stretched and the other member has been configured to the same site.

User response: Identify a node in a different site to the existing node and resubmit.

Or, change system topology.

Note: Changing the topology will disable the DR feature.

CMMVC8338E Site parameter is not supported until the current update completes.

Explanation: Site parameter is not supported until the current update completes.

User response: A node cannot be assigned to a site until the current update complete. Add the node without the -site parameter and configure sites after the update has completed.

CMMVC8339E Not supported on this system.

Explanation: Not supported on this system.

User response: The feature is not supported on this system. Wait for a later release.

CMMVC8340E Cannot modify site because system topology is stretched.

Explanation: Cannot modify site because system topology is stretched.

User response: Set topology to standard then manipulate the site.

Note: DR feature is not available when topology is standard.

CMMVC8341E Site value is not valid. Can only specify site 1 or site 2.

Explanation: Site value is not valid. Can only specify site 1 or site 2.

User response: Specify either site 1 or site 2.

CMMVC8342E Cannot set stretched topology because some nodes do not have a configured site.

Explanation: Cannot set stretched topology because some nodes do not have a configured site.

User response: Configure site for each node and then set topology.

CMMVC8343E Cannot set stretched topology because some I/O groups have 2 nodes in the same site.

Explanation: Cannot set stretched topology because some I/O groups have 2 nodes in the same site.

User response: Assign each node of the I/O group to a different site then set topology.

CMMVC8344E Cannot change site because controller has one or more managed MDisks and the system topology is stretched or Hyperswap.

Explanation: In stretched or HyperSwap system mode, it is not possible to change the site while the controller has one or more managed MDisks.

User response: Migrate or delete MDisks on the controller to make all MDisks unmanaged. Or, set the system topology to standard and then change site.

Note: The disaster recovery feature is not available with standard topology.

CMMVC8345E Cannot change site because the controller is a SAS RAID Controller.

Explanation: SAS RAID Controllers are not supported at this time.

User response: Use a FC connectivity controller.

CMMVC8346E Cannot change quorum disk because there is already a quorum disk defined at that site, and the cluster topology is stretched or hyperswap.

Explanation: Only one quorum disk is supported per site.

User response: Change either the identified current quorum disk ID, or a new quorum MDisk, so that they are not both in the same site.

Or, set the topology to standard.

Note: Setting the topology to standard will disable the DR feature.

CMMVC8347E Cannot modify rcauthmethod due to the presence of unstopped IP partnerships.

Explanation: All IP partnerships must be stopped before rcauthmethod can be changed.

User response: Use chpartnership -stop to stop partnerships, then run the command again.

CMMVC8348E Cannot modify chapsecret due to rcauthmethod set to chap and the presence of unstopped IP partnerships.

Explanation: All IP partnerships must be stopped before rcauthmethod can be changed.

User response: Use chpartnership -stop to stop partnerships, then run the command again.

CMMVC8349E The rc authentication method specified is not valid.

Explanation: Value of rcauthmethod supplied is not 'chap' or 'none'.

User response: Correct the value of the parameter.

CMMVC8350E The command failed because the specified battery is offline.

Explanation: A command to adjust battery state cannot be completed because the specified battery is offline or has been removed.

User response: If the battery is offline but present, it is safe to remove without issuing a chnodebattery -remove command. To turn on the LED, the battery must be replaced and brought online.

CMMVC8351E The command has failed because the specified battery is not redundant.

Explanation: A request to prepare a battery for removal cannot be completed because the specified battery is not redundant.

User response: Remove the condition that is causing the lack of redundancy. Such conditions may include the partner battery being offline or not fully charged, or one of the boot drives being offline.

CMMVC8352E The task cannot be initiated because the download type is not valid.

Explanation: If -type is specified, only firmware or fpga is supported now.

User response: Check the input download type and repeat the command with a supported download type.

CMMVC8353E CHAP authentication failure

Explanation: The Partner Discovery has refused a discovery request because the CHAP secret that was specified is not correct.

User response: Correct CHAP secret must be provided.

CMMVC8354E Unreachable cluster IP address

Explanation: Incorrect partner cluster IP address parameter specified.

User response: Correct IP address of the cluster must be provided. IP address must be pingable.

CMMVC8355E Remote Copy port groups not configured or incorrectly configured.

Explanation: This error occurs when the administrator has not set up ethernet ports on any node of the local system with either remote copy port group 1 or 2. In addition, they may have been failed over or are offline.

User response: The administrator must execute the cfgportip CLI to configure one or more IP addresses on any one of the local nodes with remote copy port groups 1 or 2.

CMMVC8356E Remote Copy port groups not configured or incorrectly configured.

Explanation: This error occurs when the administrator attempts to set up more than one partnership of type IPv4 or IPv6.

User response: No action possible. The only option is to remove the existing partnership and create a new partnership.

CMMVC8357E Maximum number of allowed partnerships exceeded.

Explanation: This error occurs when the administrator attempts to set up more than 3 partnerships. A maximum of 3 FC partnerships, or 2 FC and 1 IP partnerships, can exist.

User response: No action possible. The only option is to remove one of the existing partnerships and create a new partnership.

CMMVC8358E No local cluster IPs for the partnership type configured.

Explanation: This error occurs when the administrator tries to create a partnership of type IPv4, but has not configured any IPv4 type cluster IPs on the local cluster. The same error is seen if the administrator tries to create a partnership of type IPv6, but has not configured any IPv6 type cluster IPs on the local cluster.

User response: The administrator must execute the cfgportip CLI to configure the local IP address depending on the type of IP partnership which is to be created.

CMMVC8359E Partner already exists in candidate list. Cannot create partnership.

Explanation: This error is seen when an attempt is made to create an IP partnership with a remote cluster that already appears in the candidate list displayed by the lspartnershipcandidate CLI.

User response: In this scenario, an FC link exists between the clusters. The administrator must run CLI mkfcpartnership to create partnership.

CMMVC8360E Specified partner cluster IP address is used on local cluster. Cannot create partnership.

Explanation: This error is seen when the administrator specifies a local IP address as the remote cluster IP address.

User response: A local IP address cannot be used to create an IP partnership. A remote IP address must be specified.

CMMVC8361E All IP addresses of the partnership type are either down or not configured.

Explanation: This error occurs when partner discovery is not reporting any matching remote ports. Example: The partnership type is IPv4, but all ethernet ports on the partner cluster are configured with IPv6 types (or vice versa). This can also happen when the partnership type is IPv4, but all IPv4 addresses configured on the partner cluster are offline.

User response: The administrator must specify the correct partnership type. This depends on the type of IP addresses configured on the remote cluster.

CMMVC8362E The action failed because the cluster ID is not valid.

Explanation: Partnership creation failed because the specified remote cluster ID is not valid.

User response: In case of CLI `mkfcpartnership`, the administrator should run CLI `lspartnershipcandidate` to check correct cluster id/name. In case of CLI `mkippartnership`, the administrator only specifies the IP address and clusterid is discovered. If this error occurs for `mkippartnership`, contact support.

CMMVC8363E The remote cluster partnership was not created because it already exists.

Explanation: This error occurs when an attempt is made to create a partnership with a cluster that is already in a partnership.

User response: No action. A partnership cannot be created with a cluster that is already in a partnership.

CMMVC8364E Unsupported partnership type specified.

Explanation: This error occurs when the administrator specifies an unsupported type of partnership type for CLI `chpartnership`.

User response: The administrator must look for a possible option in CLI help and specify the partnership type accordingly.

CMMVC8365E Cannot change parameter if partnership is not in stopped state.

Explanation: This error occurs when the administrator tries to modify partnership parameters without stopping the partnership.

User response: Administrator must run command `chpartnership -stop clusterid/name`

CMMVC8366E Incorrect remote cluster IP specified.

Explanation: This error occurs when the administrator running `chpartnership` specifies the type as IPv4, but specifies an IPv6 address value (or vice versa).

User response: Specify an IP address that is valid for the partnership type.

CMMVC8367E Incorrect operation of FC partnership.

Explanation: While running `chpartnership`, the CLI administrator specifies option `-clusterip`, `-chapsecret`, or `-nochapsecret` for FC partnership.

User response: No action. Options `-clusterip`, `-chapsecret`, and `-nochapsecret` are not valid for FC partnership

CMMVC8369E The action failed because the cluster ID is not valid.

Explanation: User has entered wrong cluster ID.

User response: User must enter correct cluster ID.

CMMVC8370E The remote cluster partnership was not created because it already exists.

Explanation: An attempt was made to create a partnership with a cluster that is already in a partnership.

User response: No action. A partnership cannot be created with a cluster that is already in a partnership.

CMMVC8371E Unable to setup partnership due to partnership type mismatch.

Explanation: This error occurs when the partnership types of local and remote systems don't match. Example: The partnership type specified on the remote system `mkippartnership` CLI invocation is done with IPv6, but the corresponding invocation on the local system is done with IPv4.

User response: Specify the same partnership type on both clusters.

CMMVC8372E Unable to set up Remote Copy data paths with partner due to unavailability of matching Remote Copy port groups.

Explanation: This error occurs when the partner returns Remote Copy configuration information that doesn't have matching Remote Copy port group IDs. This error could also occur in cases where matching local or remote ports are offline (link state is inactive), failed over to the partner node. Example, The local system has Remote Copy port group ID set to 1 while the partner has Remote Copy port group ID set to 2.

User response: Both clusters should have the same Remote Copy group port ID.

CMMVC8373E Easy Tier is active without a license for each enclosure.

Explanation: Easy Tier must not be active without a license for each enclosure.

User response: Ensure that a license exists for each enclosure.

CMMVC8374E The command failed because the source volume has pinned data.

Explanation: A FlashCopy mapping can only be started if there is not any pinned data on the source volume.

User response: Investigate why the source volume is offline. Fix the error and bring the volume back online. Try the command again.

CMMVC8375E The command failed because the volume has pinned data.

Explanation: The command cannot be completed because of pinned data on the volume.

User response: Investigate why the volume is offline. Fix the error and bring the volume back online. Try the command again.

CMMVC8376E Cannot change quorum disk because there is already a quorum disk defined at that site, and the cluster topology is stretched or hyperswap.

Explanation: In the stretched or Hyperswap system, every quorum disk should have a valid site to ensure there is only one quorum per site.

User response: Find an MDisk with a valid site without any other quorum disks existing in this site.

CMMVC8377E Cannot change the site for a quorum drive.

Explanation: In stretched system mode, drives as a quorum disk without a valid site are not supported.

User response: Find an MDisk with a valid site without any other quorum disks existing in this site.

CMMVC8378E Cannot modify cluster IP due to the presence of unstopped IP partnerships.

Explanation: This error will display when the administrator is trying to change a clusterip and there is an active IP partnership on the cluster.

User response: Once the Administrator stops the IP partnership, the cluster IP can be changed.

CMMVC8379E Partner state is stopped.

Explanation: This error occurs when creating an IP partnership with a remote cluster and the remote cluster partnership is partially_configured_local_stopped.

User response: Run `chpartnership -start <Cluster ID>` on the remote cluster.

CMMVC8380E Partner software version mismatch.

Explanation: This error occurs when trying to create a partnership with an incompatible SAN Volume Controller software version.

User response: None. A compatible SAN Volume Controller software version is required.

CMMVC8381E The task cannot be cancelled as it has completed all drive downloads.

Explanation: After the multiple drive download task has completed, there is a 270 second delay. It is not necessary to cancel the task if it was performed within this period of time.

User response: Wait 270 seconds before trying the command again.

CMMVC8382E The system cannot apply the task because a delay has been imposed between commands. (About 270 seconds)

Explanation: When a previous `applydrivesoftware` task is completed, a 270 second delay is imposed between each command. This error message will appear if a new task is attempted during this waiting period.

User response: Wait 270 seconds before trying the command again.

CMMVC8383E System layer mismatch.

Explanation: This error occurs when creating an IP partnership between two clusters when both clusters in a different layer.

User response: The administrator should change the cluster layer so that they are the same. To do this, use command - `svctask chsystem -layer<storage/replication>`.

CMMVC8384E The command has failed because the specified fan module is offline.

Explanation: A command to adjust the fan module state can not be completed because the fan module specified is offline or has been removed.

User response: The offline fan module triggers a fan module replacement dmp. This procedure should be followed to replace the broken or missing fan module with a functional unit.

Note: Will be returned by `chenclosurefanmodule` if used on a fan module that has a current state of offline.

CMMVC8386E Sync operation not possible.

Explanation: Node boot drive sync operation is not possible.

User response: Check the `can_sync` field on the `lsnodebootdrive` view.

`can_sync` will be false if the drives are already in sync or when there are certain boot drive errors active (unsupported drive, wrong node, invalid drive contents).

In the event that drives are already in sync, do nothing.

If there are boot drive errors active, resolve the drive error before attempting to sync.

CMMVC8389E [-size] is not a supported parameter for the specified mdiskgrp.

Explanation: Cannot change the size of a parent pool.

User response: Only the `chmdiskgrp -size` command can be used to change the size of a child pool.

CMMVC8390E Cannot set *easy_tier_option* for child storage pool, the value must be same as parent storage pool's easy tier setting.

Explanation: You cannot change the Easy Tier setting of a child storage pool.

User response: If you specified the wrong storage pool, retry the command using the correct storage pool.

If you want to change the Easy Tier setting of a parent pool, use the `chmdiskgrp -easytier easy_tier_option` command.

CMMVC8412E Cannot add mdisks to a child pool.

Explanation: The `mdisk_group_id/_name` specified in `addmisk` command must be a parent pool.

This error is reported from the CLI `addmisk` command.

User response: Specify a parent pool in the `addmisk` command.

CMMVC8415E The command cannot be initiated because the maximum supported number of drives already exists.

Explanation: The system cannot manage more than 4096 drives.

User response: Remove any unused drives and try again.

CMMVC8424E The command failed because the source or target storage pool is a child storage pool, and source and target are in different parent storage pools.

Explanation: An attempt was made to run a `migratevdisk` command that specified an invalid target `mdiskgrp` (child pool) parameter. Source and target storage pools must have the same parent storage pool.

User response: Retry the command and specify a valid target.

CMMVC8427E The command failed because the child storage pool can not be shrunk below its used capacity.

Explanation: The command failed because the child storage pool cannot shrink below its used capacity.

This error is reported from `chmdiskgrp -size new_size`.

User response: Specify a new size that is larger than the used capacity of the storage pool.

CMMVC8469E The node cannot be added because it does not meet minimum hardware requirements.

Explanation: An attempt was made to add a node to the system that does not have enough cache RAM (memory) installed to run the current level of code.

User response: Either choose a different node to add, or upgrade the cache RAM (memory) on this node.

CMMVC8470E The system update cannot be completed because it is not required.

Explanation: The user is attempting to complete a system update. The action is not permitted because the system is not in a state where this is required.

User response: None - the command is not appropriate for the system.

CMMVC8471E The system update cannot be resumed because it is not stalled.

Explanation: The user is trying to resume a system update. The system update is not in the 'stalled' state so cannot be resumed.

User response: No further action is required.

CMMVC8472E The action cannot be performed because paced update is not in progress.

Explanation: The user is requesting that the next node in the system be updated as part of a paced system update. The system is not performing a paced update.

User response: No further action is required.

CMMVC8473E The next node cannot be updated because it is not ready for update.

Explanation: The user is requesting that the next node in the system be updated as part of a paced system update. The system has not finished updating the current node yet and is not ready to start updating another node.

User response: Wait until the current node update has completed and the next node is ready.

CMMVC8474E The node cannot be updated because it would cause VDisks to go offline.

Explanation: You are requesting that the next node in the system be updated as part of a paced system update. If this node is updated then some volumes will temporarily go offline.

User response: Either address the issue causing the volumes not to be redundant, or resubmit the command using the `-force` option. If the `-force` option is used then some volumes will temporarily go offline.

CMMVC8475E The node cannot be updated because it is offline.

Explanation: You are requesting that the next node in the system be updated as part of a paced system update. The node cannot be updated because it is offline.

User response: Either delete the node from the system or perform service actions to bring it back online.

CMMVC8476E The node cannot be added because it would cause a paced update but the system has not completed its current update.

Explanation: You are trying to add a node that has

been configured to perform a paced update. A system update is currently in progress and it is not a suitable time to begin a paced update.

User response: Wait until the system update has completed, then add the node again.

CMMVC8477E The command failed because firmware applicable to a Storwize V7000 Gen2 drive was not found in the package.

Explanation: While updating drives on the enclosure, the CLI returns this error if an incorrect format was used for the file name of the drive update package.

The correct file name format is:

`mdisksw.product_id.firmware.fw_level.img`

User response: Check the .img file name in the .gpf file and retry the command.

CMMVC8478E The mapping cannot be removed because the volume being unmapped has received I/O within the defined volume protection period.

Explanation: Volume protection is enabled, but the volume being unmapped has received I/O recently, within the defined volume protection time period, so the unmapping command fails. Receiving I/O generally indicates that a volume is in use.

User response:

1. Ensure that this volume is intended to be unmapped. If you selected the wrong volume, repeat the unmapping command with the correct volume.
2. To unmap this volume, ensure that no host I/O is sent to the volume, wait the time specified in the `vdisk_protection_time` field in the `lssystem` command since the last I/O was received, then retry the unmapping command.
3. To disable volume protection and its warnings, disable the `vdisk_protection-enabled` field in the `chsystem` command.

CMMVC8479E The volume cannot be removed because the volume has received I/O within the defined volume protection period.

Explanation: If volume protection is enabled and an attempt to delete the last volume copy occurs within the defined volume protection time period after the volume received I/O, the volume removal command fails. This policing behavior protects I/O integrity.

This policing behavior:

- Is not affected by the volume being mapped or unmapped

- Does not occur when deleting one of two volume copies
- Is not overridden by using the **force** flag

User response:

1. Ensure that you intend to unmap the volume. If you selected the wrong volume, repeat the command with the correct volume.
2. To unmap the volume, ensure that no host I/O is sent to the volume, wait the time that is specified in the `vdisk_protection_time` field in the **lssystem** command, then retry the volume removal command.
3. Alternatively, to disable volume protection warnings and behavior, disable the `vdisk_protection-enabled` field using the **chsystem** command, then retry the volume removal command.

CMMVC8481E The storage pool cannot be deleted because at least one volume in the pool has received I/O within the defined volume protection period.

Explanation: If volume protection is enabled, and any volume in the storage pool being deleted received I/O within the defined volume protection time period, the storage pool deletion command fails. This is a policing behavior to protect I/O integrity.

This policing behavior is not overridden by the **force** flag.

User response:

1. Ensure that no volumes are in use in the pool. If you selected the wrong pool, repeat the command with the correct pool.
2. To unmap all volumes in the pool, ensure that no host I/O is sent to any volume in the pool, wait the time that is specified in the `vdisk_protection_time` field in the **lssystem** command, then retry the storage pool deletion command.
3. Alternatively, to disable volume protection warnings and behavior, disable the `vdisk_protection-enabled` field using the **chsystem** command, then retry the storage pool deletion command.

CMMVC8482E The host cannot be removed because the host being deleted is mapped to volumes which have received I/O within the defined volume protection period.

Explanation: If volume protection is enabled, and the host being deleted is mapped to any volume that has received I/O within the defined volume protection time period, the host deletion command fails. This policing behavior protects I/O integrity.

This policing behavior:

- Does not occur when multiple hosts are mapped to the same volume, and the host being removed is already offline

- Is not overridden by using the **force** flag
- Prevents the last host being removed regardless of whether the host is online or not, which improves the policing behaviour for clustered hosts

User response:

1. Ensure that you intend to remove the host. If you selected the wrong host, repeat the command with the correct host.
2. To remove the host, ensure that no host I/O is sent to any volume mapped to the host, wait the time that is specified in the `vdisk_protection_time` field in the **lssystem** command, then retry the volume removal command.
3. Alternatively, to disable volume protection warnings and behavior, disable the `vdisk_protection-enabled` field using the **chsystem** command, then retry the host removal command.

CMMVC8483E The host I/O group cannot be removed because at least one volume in the I/O group(s) being removed from the host has received I/O within the defined volume protection period.

Explanation: If volume protection is enabled, but any volume mapped to the host in the removal candidate I/O group has received I/O recently (within the defined volume protection time period), the host I/O group removal command, **rmhostiogrp**, fails.

This behavior is not affected by the **force** parameter of the command. The force parameter deletes from a host the I/O groups that have volume mappings. The force parameter does not override the volume protection policing.

Removing a host I/O group behaves the same way as does removing a host in that the policing behavior allows the removal of an I/O group if the host is offline, and if one other host is mapped to the volume. The policing policy does not allow removing the I/O group if the host is the only host mapped to the volume and the volume is busy.

User response:

1. Ensure that the host I/O group is intended to be unmapped. If you selected the wrong host I/O group, repeat the removal command with the correct host I/O group.
 2. To unmap the volumes in the host I/O group and remove the I/O group, ensure that no host I/O is sent to the volumes, wait the time specified in the `vdisk_protection_time` field in the **lssystem** command since the last I/O was received, then retry the host I/O group removal command.
 3. To disable volume protection and its warnings, disable the `vdisk_protection-enabled` field in the **chsystem** command.
-

CMMVC8484E The port cannot be removed because the host is mapped to at least one volume which have received I/O within the defined volume protection period and the port is the last port associated with the host.

Explanation: If volume protection is enabled, but the host port is the last port for a host mapped to any volume that has received I/O recently (within the defined volume protection time period), the **rmhostport** command fails.

The **rmhostport** command shares a similar behavior with the **rmhost** and the **rmhostiogrp** commands: volume protection policing allows removal of the last host port if the host is already offline and there is another host mapped to the volume. The policing policy does not allow removing the host port if the host is the only host mapped to the volume and the volume is busy.

This behavior is not affected by the **force** parameter of the command.

User response:

1. Ensure that the host port is intended to be removed. If you selected the wrong host port, repeat the removal command with the correct host port.
2. To unmap the volumes in the host I/O group and remove the host port, ensure that no host I/O is sent to the volumes, wait the time specified in the **vdisk_protection_time** field in the **lssystem** command since the last I/O was received, then retry the host I/O group removal command.
3. To disable volume protection and its warnings, disable the **vdisk_protection-enabled** field in the **chsystem** command.

CMMVC8485E The command cannot be initiated because the drive that you have specified has not been formatted with the correct protection information scheme.

Explanation: The drive has not been formatted to the correct protection information scheme (Type 2), so initialization cannot be performed. This usually implies a process step has been missed during the manufacturing of the drive.

User response: If you are confident that you have no data on the drive, use **chdrive -task format <drive_id>** to reformat the drive to the correct protection information scheme. Then, retry the command.

CMMVC8518E The command failed because image-mode volumes are not supported in child storage pools.

Explanation: You cannot create an image mode volume from the child pools.

This error is reported by the **mkvdisk -mode image** command.

User response: Change the volume mode to **striped**.

CMMVC8519E The command failed because sequential-mode volumes are not supported in child storage pools.

Explanation: You cannot create a sequential mode volume from the child pools.

This error is reported by the **mkvdisk -mode seq** command.

User response: Change the volume mode to **striped**.

CMMVC8523E The command failed because the specified storage pool is a child storage pool.

Explanation: This error is reported by the **mkmdiskgrp -parentmdiskgrp** command. The **parentmdiskgrp** parameter must identify a parent storage pool. You cannot create a child pool from another child pool.

User response: Specify a parent storage pool for the **-parentmdiskgrp** parameter.

CMMVC8525E The command failed because mdisks cannot be added to a child storage pool.

Explanation: The **mdisk_group_id** or **mdisk_group_name** that was specified in the **addmdisk** command must be a parent pool.

User response: Specify a parent pool in the **addmdisk** command.

CMMVC8526E The command failed because mdisks cannot be removed from a child storage pool.

Explanation: The **mdisk_group_id** or **mdisk_group_name** that was specified in the **rmmdisk** command must be a parent pool.

User response: Specify a parent pool in the **rmmdisk** command.

CMMVC8528E The system does not support encryption.

Explanation: The hardware on this system does not support encryption.

User response: Use a system hardware that supports encryption.

CMMVC8529E The encryption function has not been activated.

Explanation: The system supports encryption but licenses have not been activated.

User response: Activate the encryption function on all I/O groups by installing license keys.

CMMVC8530E A re-key operation is in progress.

Explanation: Command can't be carried out whilst rekey operation is in progress.

User response: Run chencryption to commit or cancel current rekey operation.

CMMVC8531E Encryption cannot be disabled while functions are configured to use encryption.

Explanation: The administrator tried to disable encryption when there are still encryption keys in use and there is presumed to be encrypted data somewhere on the system.

User response: Remove all encrypted objects and try again, or leave encryption enabled while encryption is in use.

CMMVC8532E The encryption function is not enabled.

Explanation: The encryption key cannot be set up until encryption is enabled.

User response: If your system supports encryption, enable encryption and try again. You can find more information by searching for "Enabling encryption".

CMMVC8533E No re-key operation in progress.

Explanation: No re-key operation to commit or cancel.

User response: Prepare key first and retry commit task.

No action required to cancel.

CMMVC8534E node [X] has insufficient entropy to generate key material.

Explanation: An internal error in encryption has occurred.

User response: The node used for key generation needs to be rebooted or replaced.

CMMVC8535E The system does not have access to current encryption keys.

Explanation: The system does not have access to the USB flash drives containing current encryption keys.

User response: Locate the USB flash drive containing encryption keys and plug them into the system.

CMMVC8536E The system does not have sufficient USB flash drives attached.

Explanation: The system does not have sufficient USB flash drives attached. The number needed is displayed.

User response: Insert the specified number of USB flash drives needed into the system and retry.

CMMVC8537E Not enough USB copies made, needs (%1).

Explanation: The keys have not been written to the minimum number of USB devices.

User response:

- Wait for the copies count to reach the minimum required
- Check for USB errors
- Replace flash drive and cancel. Then, retry prepare.

CMMVC8538E System is busy commit in progress.

Explanation: The system is busy. Commit in progress.

User response: Wait for I/O to complete and commit to complete. Then, use lsencryption to check state.

CMMVC8539E Encryption is not supported on drives IO group

Explanation: The drive selected is in an I/O group that does not support encryption.

User response: Select a drive in an I/O group that supports encryption.

CMMVC8540E Key not accepted by SAS adapter because of an internal error.

Explanation: Key not accepted by SAS adapter.

User response: Follow DMP for error shown in lsencryption.

CMMVC8541E The command cannot be initiated because the object is being used by an automatic system configuration task.

Explanation: This message appears when:

- A failed drive has been reseated or replaced. The system has automatically configured the device. A drive was automatically managed in the configuration as a replacement for a failed drive. This might have been due to user action of replacing the failed hardware with a new drive or reseating the old hardware on advice from a DMP.

- An attempt to automatically configure a reseated or replaced drive failed.
A failed drive has been reseated or replaced. The system attempted to automatically configure the reseated or replaced drive for use in the system. This attempt failed.

User response: Wait until the drive auto manage process is complete. An informational or error message appears in the event log.

CMMVC8542E The system update test utility has not been run before updating the system.

Explanation: Before you can update your system, you need to run the latest version of the test utility to verify that there are no issues with the current system.

User response: Run the system update test utility before you attempt a system update.

CMMVC8544E The command cannot be initiated because the object is being used by an automatic system configuration task.

Explanation: The chenclosureslot command will no longer work while this new drive auto manage code is running.

User response: Wait until drive automanage finishes before running the chenclosureslot command.

CMMVC8549E The command failed because the source volume is from a child pool.

Explanation: The volumes of a child pool can be migrated only to its parent pool or to a child pool that comes from the same parent.

User response: Choose a target pool that has the same parent as the source volume and retry the command.

CMMVC8550E Cluster ID alias cannot be changed whilst encryption is enabled.

Explanation: Encryption requires that the cluster ID alias does not change while encryption is enabled.

User response: To change the cluster ID alias, disable encryption.

Once completed, re-enable encryption.

CMMVC8567E A host of type adminlun is required to be mapped to all IO groups.

Explanation: An attempt was made to change the host type to adminlun, but the host is not in all I/O groups.

User response: Add the host to all I/O groups by using the following command:
`addhostiogrp -iogrpall host_id`

You can now change the host type to adminlun.

CMMVC8570E A subsidiary volume may only be mapped to hosts of type adminlun.

Explanation: An attempt was made to map a subsidiary volume to a host that wasn't an adminlun type. Subsidiary LUN mappings are typically created or removed by the IBM Spectrum Control Base Edition application. These mappings apply only to adminlun host types, and so it is unlikely that a storage administrator would encounter this message.

User response: Contact IBM Support.

CMMVC8571E The volume is already mapped as a subsidiary volume.

Explanation: An attempt was made to map a volume as a subsidiary volume, but the volume is already mapped as one. Subsidiary LUN mappings are typically created or removed by the IBM Spectrum Control Base Edition application. These mappings apply only to adminlun host types, and so it is unlikely that a storage administrator would encounter this message.

User response: Contact IBM Support.

CMMVC8573E The action failed because there is one special volume owned by metadata volume in this storage pool.

Explanation: The `rmmdiskgrp` command failed even though the `-force` parameter was specified. One special volume in the specified storage pool is owned by the metadata volume.

User response: Complete the following steps:

1. Run the `lsmetadatavdisk` command to find the special volume that is owned by the metadata volume.
2. Run the `migratevdisk` command to migrate the special volume from the specified storage pool to a different storage pool.
3. Retry the `rmmdiskgrp` command.

CMMVC8574E The action failed because some volumes with owner type vvol depend on the metadatavdisk and -ignorevvolsexist was not specified.

Explanation: An attempt was made to remove a metadata volume that some other volumes depend on.

User response: Complete one of the following actions:

- Remove the vvol type volumes by using the `rmvdisk` command, then retry the `rmmetadatavdisk` command.
- Alternatively, specify the `-ignorevvolsexist` parameter when you retry the command. Using this

parameter means that the metadata volume is removed, but the dependent vvol volumes are retained.

CMMVC8575E The action failed because it depends on NTP server but NTP server is not configured.

Explanation: During the creation of the metadata volume, an incorrect timestamp was found. The error was caused by the lack of a network time protocol (NTP) server, which is required for the metadata volume.

User response: Create the NTP server by using the following command:

```
chsystem -ntpip
```

After the NTP server is created, retry the **mkmetadatavdisk** command.

CMMVC8590E The command failed because the Compare Volume ID or name is invalid.

Explanation: The specified Compare Volume ID or name is not valid.

User response: Specify a valid Compare Volume ID or name and resubmit the command.

CMMVC8591E The command failed because the Base Volume ID or name is invalid.

Explanation: The command failed because the Base Volume ID or name is invalid.

User response: Specify a Base Volume ID or name and resubmit the command.

CMMVC8592E The command failed because the Compare Volume does not exist.

Explanation: The command failed because the Compare Volume does not exist.

User response: Specify the ID or name of an existing Compare Volume and resubmit the command.

CMMVC8593E The command failed because the Base Volume does not exist.

Explanation: The command failed because the Base Volume does not exist.

User response: Specify the ID or name of an existing Base Volume and resubmit the command.

CMMVC8594E The command failed because the Base and Compare Volumes are the same.

Explanation: The command failed because the Base

and Compare Volumes are the same.

User response: Specify Base and Compare Volumes that are not identical and resubmit the command.

CMMVC8595E The command failed because the startlba of the scan segment is invalid.

Explanation: The command failed because the startlba of the scan segment is invalid.

User response: Specify a valid startlba and resubmit the command.

CMMVC8596E The command failed because the lbacount of the scan segment is invalid.

Explanation: The command failed because the lbacount of the scan segment is invalid.

User response: Specify a valid lbacount and resubmit the command.

CMMVC8597E The command failed because the chunk size is invalid.

Explanation: The command failed because the specified chunk size is not valid.

User response: Specify a valid chunk size and resubmit the command.

CMMVC8598E The command failed because the volume and the base volume are not in the same dependency chain.

Explanation: The command failed because the volume and the base volume are not in the same dependency chain.

User response: Specify base and compare volumes that are in the same dependency chain and resubmit the command.

CMMVC8599E The command failed because a node is offline.

Explanation: A node is offline which caused the command to fail.

User response: Bring the node back online and resubmit the command.

CMMVC8600E The command failed because the write cache could not be flushed in time.

Explanation: The command failed because the write cache could not be flushed in time.

User response: Retry the command.

CMMVC8601E The command failed because the start of the scan segment is not chunk aligned.

Explanation: The command failed because the start of the scan segment is not chunk-aligned.

User response: Specify a startlba that is chunk-aligned and resubmit the command.

CMMVC8620E The active-active relationship could not be created because the master and auxiliary volumes do not have different well-defined sites.

Explanation: The active-active relationship could not be created because the master and auxiliary volumes do not have different well-defined sites.

User response: Contact IBM Support.

CMMVC8621E The active-active relationship could not be created because the system topology is not HyperSwap.

Explanation: The active-active relationship could not be created because the system topology is not HyperSwap.

User response: Contact IBM Support.

CMMVC8622E The active-active relationship could not be created because the master and auxiliary volumes are not in the same system.

Explanation: The active-active relationship could not be created because the master and auxiliary volumes are not in the same system.

User response: Contact IBM Support.

CMMVC8623E I/O group was not specified. I/O group must be specified because topology is HyperSwap.

Explanation: You cannot set HyperSwap topology without specifying the I/O group.

User response: Specify the I/O group when using HyperSwap technology.

CMMVC8624E Cannot set HyperSwap topology because some nodes do not have a configured site.

Explanation: Cannot set HyperSwap topology because some nodes do not have a configured site.

User response: Configure the site for each node and then set the topology.

CMMVC8625E Cannot set HyperSwap topology because an I/O group has nodes assigned to different sites.

Explanation: You cannot set HyperSwap topology when the I/O group has nodes assigned to different sites.

User response: Assign both nodes in the I/O group to the same site or do no attempt to set HyperSwap.

CMMVC8626E Cannot modify site because topology is HyperSwap.

Explanation: You cannot change the site of a node because the topology is set to HyperSwap.

User response: Change the system topology to standard to change the site of the node.

Note: Changing the topology disables the disaster recovery (DR) feature.

CMMVC8627E Site was not specified. Site must be specified because topology is Hyperswap.

Explanation: Site was not specified. Site must be specified because topology is HyperSwap.

User response: Identify the site and resubmit command with -site flag. Alternatively change system topology.

Note: Changing the topology disables the disaster recovery feature.

CMMVC8628E The host mapping was not created because the volume is not the host accessible volume in an active-active relationship.

Explanation: A volume cannot be mapped to a host if it is an auxiliary volume in an active-active relationship.

User response: Create the host mapping to the master volume in the relationship.

CMMVC8629E Site specified is not valid. System topology is Hyperswap and the other member of the IO group has been configured to a different site.

Explanation: Site specified is not valid. System topology is HyperSwap and the other member of the I/O group has been configured to a different site.

User response: Identify the site of the other member of the I/O group. Assign the same site and resubmit. Alternatively, change the system topology.

Note: Changing the topology will disable the DR feature.

CMMVC8630E Site specified is not valid. System topology is HyperSwap and the node or control enclosure was previously configured with a different site.

Explanation: Site specified is not valid. System topology is HyperSwap and the node or control enclosure was previously configured with a different site.

User response: Identify the site specified originally for this node or control enclosure and resubmit the command. Or, change the system topology.

Note: Changing the topology disables the disaster recovery feature.

CMMVC8631E The change VDisk could not be disassociated because it is configured in an active-active relationship.

Explanation: A change volume cannot be disassociated from an active-active relationship.

User response: Contact IBM Support.

CMMVC8632E The copy type of an active-active relationship or group cannot be changed.

Explanation: The copy type of an active-active relationship or group cannot be changed.

User response: Contact IBM Support.

CMMVC8633E The relationship could not be added to the consistency group because it does not have an up-to-date copy in the same site as a set of up-to-date copies in the group.

Explanation: When adding an active-active relationship to a consistency group, the relationship must have an up-to-date copy in the same site as a set of up-to-date copies of each relationship in the consistency group.

User response: Before attempting to add the relationship to the consistency group, verify that the copy in the candidate relationship is as up-to-date as the set of up-to-date copies of any existing relationship in the consistency group.

CMMVC8634E The host mapping was not created because the host does not have a site defined and the volume is in an active-active relationship.

Explanation: A volume in an active-active relationship

can only be mapped to a host if the host has a site defined.

User response: Set the site for the host using the management GUI or the lshost and chhost CLI commands.

CMMVC8635E Cannot un-set the host site as the host is mapped to a volume in an active-active relationship.

Explanation: A host mapped to a volume in an active-active relationship must have a site defined.

User response: Keep the site defined or remove the host mapping if un-setting the host site is required.

CMMVC8636E The active-active relationship was not created because one or more hosts mapped to the master volume do not have a site defined.

Explanation: A host mapped to a volume in an active-active relationship must have a site defined.

User response: Set the site of all hosts that are mapped to the master volume before resubmitting the command.

CMMVC8637E The active-active relationship could not be created because there are one or more host mappings to the auxiliary volume.

Explanation: The active-active relationship could not be created because there are one or more host mappings to the auxiliary volume.

User response: Contact IBM Support.

CMMVC8638E An active-active relationship or group may not be manually started unless it has a state of idling and access has previously been enabled to an out-of-sync consistent copy.

Explanation: An active-active relationship or group may not be manually started unless it has a state of idling and access has previously been enabled to an out-of-sync consistent copy.

User response: Contact IBM Support.

CMMVC8639E An active-active relationship or group may not be manually stopped unless it has a state of consistent_copying and access is required to the out-of-sync consistent copy.

Explanation: An active-active relationship or group may not be manually stopped unless it has a state of consistent_copying and access is required to the out-of-sync consistent copy.

User response: Contact IBM Support.

CMMVC8640E Active-active relationships and groups alter their direction automatically and cannot have their direction switched manually.

Explanation: Active-active relationships and groups alter their direction automatically and cannot have their direction switched manually.

User response: Contact IBM Support.

CMMVC8642E The command failed because the copy specified is not synchronized. The copy must be synchronized to create an active-active relationship.

Explanation: Only volumes where both copies are synchronized can be split when creating an active-active relationship. The command failed because the specified copy is not synchronized.

User response: Contact IBM Support.

CMMVC8643E The active-active relationship could not be created because the existing and new volumes do not have different well-defined sites.

Explanation: The active-active relationship could not be created because the existing and new volumes do not have different well-defined sites.

User response: Contact IBM Support.

CMMVC8644E The command failed because the auxiliary volume does not contain an up-to-date copy or have host access enabled.

Explanation: The master volume in an active-active relationship (keeping host access to the auxiliary volume) can only be removed when the auxiliary volume is already processing host I/O. Either it is an up-to-date copy of an active-active relationship or the copy has host access enabled with the "stoprelationship -access" command.

User response: Contact IBM Support.

CMMVC8646E The active-active relationship could not be created, either because the master or auxiliary volume is the target of a FlashCopy mapping or because the master or auxiliary volume is the source of a FlashCopy mapping with a target volume in a different site.

Explanation: The active-active relationship could not be created because:

- The master or auxiliary volume is the target of a FlashCopy mapping.
- The master auxiliary volume is the source of a FlashCopy mapping with a target volume in a different site.

User response: Contact Support.

CMMVC8649E Cannot change topology because active-active relationships are defined, and only supported in Hyperswap topology.

Explanation: It is not possible to change the topology because active-active relationships are defined and only supported in a HyperSwap topology.

User response: Remove all active-active relationships before changing the topology.

CMMVC8650E System topology is Hyperswap or stretched, and the site of the MDisk being added does not match that of the storage pool.

Explanation: The system topology is either HyperSwap or stretched and the site of the MDisk you are adding does not match the storage pool.

User response: Do one of the following:

- Add the MDisk to an empty storage pool or to one already containing MDisk of a matching site.
- Remove and re-add the controller of the MDisk in the desired site.

Then, retry the command.

CMMVC8652E The command failed because the volume is owned and has restricted use.

Explanation: You cannot specify this command for this volume because the volume is owned and has restricted use.

User response: Choose a volume that is not owned with restricted use.

CMMVC8653E The command failed as an MDisk group is owned and has restricted use.

Explanation: An MDisk group specified in the command is owned and either the action is not allowed for this type of owner or the user requires a specific role.

User response: Check if the MDisk group is owned and then make sure the command is permitted for this type of owner or the required role is being used.

CMMVC8654E The storage pool specified is not valid. The volume is a participant in an active-active relationship, and the storage pool for the new copy is in a different site to the volume's current site.

Explanation: Participants in an active-active relationship must be located in the same site.

User response: Move the volume to a storage pool in the same site, or move the other copy in the active-active relationship to the target storage pool.

CMMVC8655E The volume being moved is a participant in an active-active relationship, and the target storage pool is in a different site to the volume's current site.

Explanation: The members of an active-active relationship must be located on the same site.

User response: Move the volume to a storage pool in the same site, or move the other copy in the active-active relationship to the target storage pool.

CMMVC8658E The create failed because there are no iogroups supporting compression.

Explanation: An attempt was made to create a compressed volume where the **-iogrp** parameter was not specified and none of the candidate I/O groups for the command support compression.

User response: Retry the command and use the **-iogrp** parameter to specify an I/O group that supports compression. Alternatively, retry the command without the **-compressed** parameter.

CMMVC8659E The Create FlashCopy mapping task cannot be initiated because the target volume is in an active-active relationship.

Explanation: An attempt was made to create a FlashCopy mapping with a HyperSwap volume as a target. HyperSwap volumes cannot be used as targets for FlashCopy mappings. This rule applies to both the master and auxiliary volumes of the active-active relationship.

User response: If you must use FlashCopy to create a point-in-time copy of a volume onto a HyperSwap volume, that target HyperSwap volume must be reconfigured as a regular volume until the FlashCopy process completes. Use the management GUI or the **rmvolumecopy** command to delete one copy of the target HyperSwap volume before you create the FlashCopy mapping.

CMMVC8660E The Create FlashCopy mapping task cannot be initiated because the source volume is in an active-active relationship and the target volume and map are not in the same site as the source volume.

Explanation: HyperSwap volumes contain a copy on each site, and copying data from one site to the other introduces significant unnecessary data movement. When you create a FlashCopy mapping to copy data on a HyperSwap volume, the source volume of the FlashCopy mapping must be the HyperSwap volume copy that is on the same site as the target volume.

User response: Retry the command, specifying as a source volume the HyperSwap volume copy that is on the same site as the target volume.

CMMVC8661E The command has failed because the specified volume is a target of a FlashCopy mapping with a source volume in an active-active relationship and the new I/O group is in a different site to the source volume.

Explanation: An attempt was made to move the preferred node of a volume to an I/O group on a different site when the volume is a target of a FlashCopy mapping with a source volume in an active-active relationship. For FlashCopy mappings that copy data from a HyperSwap volume, the source volume of the FlashCopy mapping must be on the same site as the target volume.

User response: Retry the command and move the volume to another I/O group on the same site. Alternatively, to move the volume to an I/O group on the other site, remove the FlashCopy mapping or convert the HyperSwap source volume to a regular volume by using the management GUI or the **rmvolumecopy** command first.

CMMVC8662E The command has failed because the specified volume is a target of a FlashCopy mapping with a source volume in an active-active relationship and the target storage pool is in a different site to the source volume.

Explanation: HyperSwap volumes contain a copy on each site, and copying data from one site to the other introduces significant unnecessary data movement. For a FlashCopy mapping that copies data from a HyperSwap volume, the source volume of the FlashCopy mapping must be on the same site as the target volume.

User response: Retry the command and move the volume to a storage pool on the same site. Alternatively, to move the volume to a storage pool on the other site, convert the HyperSwap source volume to

a regular volume by using the management GUI or the **rmvolume** command first.

CMMVC8663E The command has failed because the specified volume is a target of a FlashCopy mapping with a source volume in an active-active relationship and the storage pool for the new copy is in a different site to the source volume.

Explanation: HyperSwap volumes contain a copy on each site, and copying data from one site to the other introduces significant unnecessary data movement. For a FlashCopy mapping that copies data from a HyperSwap volume, the source volume of the FlashCopy mapping must be entirely on the same site as the target volume. This command would spread the volume across both sites.

User response: Complete one of the following actions:

- Delete the FlashCopy mapping and retry the command.
- Move the FlashCopy target volume to the same site as the source volume:
 1. Use the **movevdisk** command to change the I/O group for the target volume to one on the same site as the source volume.
 2. Use the **migratevdisk** command to change the storage pool for the target volume to a storage pool on the same site as the source volume.

Then retry the command.

CMMVC8664E The change volume could not be associated because it does not have the same well-defined site as the associated volume in the relationship.

Explanation: Change volumes support HyperSwap volumes by capturing consistent data during synchronization. This process allows access to the consistent data if the remote site is lost. The change volume must be on the correct site for the volume in the active-active relationship it is being associated with.

User response: Configure the change volume in an I/O group and storage pool on the same site as the volume in the active-active relationship you are associating it with.

CMMVC8665E The FlashCopy mapping was not started because the source volume is in an active-active relationship and does not contain an up-to-date copy or have host access enabled.

Explanation: HyperSwap volumes contain a copy on each site, and copying data from one site to the other introduces significant unnecessary data movement. For a FlashCopy mapping that copies data from a HyperSwap volume, the HyperSwap volume must

have an up-to-date copy on the same site as the target volume of the FlashCopy mapping.

User response: Allow the HyperSwap volume to synchronize fully before you start the FlashCopy mapping. Alternatively, create and use a FlashCopy mapping on the same site as the up-to-date copy of the HyperSwap volume.

CMMVC8666E The FlashCopy consistency group was not started because one or more of the source volumes are in an active-active relationship and do not contain an up-to-date copy or have host access enabled.

Explanation: HyperSwap volumes contain a copy on each site, and copying data from one site to the other introduces significant unnecessary data movement. For a FlashCopy consistency group that has one or more maps that copy data from a HyperSwap volume, the HyperSwap volume must have an up-to-date copy on the same site as the target volume of the FlashCopy mapping.

User response: Allow the HyperSwap volume to synchronize fully before you start the FlashCopy mapping. Alternatively, create and use a FlashCopy mapping on the same site as the up-to-date copy of the HyperSwap volume.

CMMVC8668E The host mapping was not created because the SCSI logical unit number (LUN) ID is not valid for this particular host type.

Explanation: An attempt was made to create a mapping between a volume and an adminlun host, but the SCSI LUN ID is not within the range for adminlun hosts. The highest SCSI LUN ID that is permitted for an adminlun host is 512.

User response: Retry the command with a valid SCSI LUN ID.

CMMVC8670E Cannot modify host type since it has subsidiary mappings.

Explanation: An attempt was made to change the host type from adminlun to another host type, but the adminlun host has subsidiary mappings. Only adminlun hosts can have subsidiary mappings.

User response: Contact a VMware administrator to ensure that all Virtual Volume datastores are unmounted from the specified host because this operation removes all remaining subsidiary mappings. After the process is complete, retry the command.

CMMVC8671E The host's SCSI ID exceeds the maximum for this host type.

Explanation: An attempt was made to change the host type to adminlun, but no free SCSI IDs were available for this host type. The maximum number of SCSI IDs that are available for the adminlun host type is 512.

User response: Make a SCSI ID available for the new adminlun host by completing one of the following steps:

- Delete an adminlun host by using the **rmhost** command.
- Change the type of an existing adminlun host by using the **chhost** command. Note that if you are using virtual volumes, you cannot change the host type.

CMMVC8672E The command failed because the access IO group set of a subsidiary volume may not be modified.

Explanation: An attempt was made to modify the access I/O group set for a subsidiary volume, which is required to be available through all I/O groups.

User response: You cannot modify access of a subsidiary volume. Retry the command for a different volume.

CMMVC8673E The volume is already mapped as a SCSI LUN which must first be unmapped/removed in order to create a subsidiary volume mapping.

Explanation: An attempt was made to create a subsidiary volume mapping, but the volume is already mapped as a SCSI LUN. Subsidiary LUN mappings are typically created or removed by the IBM Spectrum Control Base Edition application. These mappings apply only to adminlun host types, and so it is unlikely that a storage administrator would encounter this message.

User response: Contact IBM Support.

CMMVC8674E The volume is not mapped as a subsidiary LUN.

Explanation: An attempt was made to remove a subsidiary mapping between an adminlun host and a volume, but the volume was not a subsidiary volume.

User response: Specify a valid volume ID and retry the command.

CMMVC8676E The command failed because the volume is not a master volume in an active-active relationship.

Explanation: This command retains access to a HyperSwap volume when deleting the master volume

of its active-active relationship. It is not necessary or permitted when deleting the auxiliary volume of its active-active relationship, or any non-HyperSwap volume.

User response: Retry the command on the correct object type.

CMMVC8677E Retaining access to the auxiliary VDisk failed because either the master or the auxiliary VDisk is in mirrored mode.

Explanation: The master volume in an active-active relationship can only be removed, keeping host access to the auxiliary volume, when both the master and auxiliary volumes are in non-mirrored volume.

User response: Remove additional mirrors of the master and auxiliary volumes of the active-active relationship so each has only one mirror, then retry the command.

CMMVC8697E The command failed because the source or target volume is in the deleting state.

Explanation: An attempt was made to delete a volume that was already in the process of being deleted.

User response: Complete the following actions:

1. Make sure that you specified the correct volume.
2. If so, wait until the previous **rmvdisk** command completes.
3. If you believe that the previous command failed, and you are sure that you want to proceed, you can retry the command with the **-force** parameter. Using the **-force** parameter bypasses all built-in safety checks.

CMMVC8698E The command failed because the volume has host mappings, use -removehostmappings.

Explanation: An attempt was made to delete a volume that has host mappings.

User response: Make sure that you specified the correct volume. If so, and you want to delete the host mappings along with the volume, retry the command and include the **-removehostmappings** parameter.

CMMVC8699E The command failed because the volume is image mode and has FlashCopy maps, either remove the FlashCopy maps first or use -force.

Explanation: An attempt was made to delete an image mode volume that has FlashCopy maps.

User response: Complete the following steps:

1. Make sure that you specified the correct volume.
2. If so, remove the FlashCopy maps by using the **rmfcmaps** command and retry the **rmvdisk** command.
3. Alternatively, you can retry the **rmvdisk** command with the **-force** parameter. Using the **-force** parameter bypasses all built-in safety checks.

CMMVC8706E The command failed because the primary copy's autodelete flag is on, meaning type conversion is in progress.

Explanation: Changing the primary copy (using the **chvdisk -primary** command) is disabled when the primary copy's autodelete flag is on.

User response: Change the primary copy when the primary copy's **autodelete** is **no**.

Either wait for the type conversion to finish or cancel the process, if required, by using the **rmvdiskcopy** command.

CMMVC8707E The command failed because the caching IO groups for the volume are not in two independent sites.

Explanation: When you create HyperSwap volumes, the caching I/O groups must be located in both site 1 and site 2.

User response: Ensure that you have caching I/O groups in both site 1 and site 2, then retry the command.

CMMVC8708E The command failed because the system does not have a hyperswap topology.

Explanation: An attempt was made to create a HyperSwap volume on a system that does not have a HyperSwap topology. The **mkvolume** command attempts to create a HyperSwap volume in either of the following circumstances:

- Two I/O groups are specified in the command
- Two storage pools are specified within two independent sites that have a HyperSwap topology. In this case the **iogrp** parameter can be omitted and the system will auto-select an I/O group.

User response: Complete one of the following actions:

- If you entered the command incorrectly, correct the command and retry.
- Enter the **mkvolume** command with a single **-iogrp** parameter to create a basic volume.
- Configure the system to use a HyperSwap topology and retry the command.

CMMVC8709E The command failed because the caching IO groups are not in the same sites as the storage pools.

Explanation: When you create volumes under a HyperSwap topology, the caching I/O groups must be in the same sites as the storage pools.

User response: Ensure that the caching I/O groups are in the same sites as the storage pools, then retry the command.

CMMVC8710E The command failed because there were not enough extents in storage pool *storage_pool*.

Explanation: Volumes are created from the extents that are available in a storage pool. An attempt was made to create a volume when not enough extents were available.

User response: If you specified the correct storage pool, complete one of the following tasks:

- Reduce the size of the volume you are creating.
- Add more objects to the pool.

Then, retry the command.

CMMVC8711E The command failed because too many IO groups were specified.

Explanation: The maximum number of I/O groups is 2.

User response: Retry the command, specifying 2 or fewer I/O groups.

CMMVC8712E The command failed because too many storage pools were specified.

Explanation: An attempt was made to create a volume or volume copy that referenced more than two storage pools. The following limits are enforced:

- For the **mkvdisk** command: two mirrored copies.
- For the **mkvolume** command with standard or stretched topology: two mirrored copies.
- For the **mkvolume** command with HyperSwap topology: two volume copies.

User response: Retry the command with no more than two storage pools specified.

CMMVC8713E The command failed because the storage pool must be in site 1 or 2.

Explanation: When you create a volume under a stretched or HyperSwap topology, the storage pool must be located in either site 1 or site 2.

User response: Ensure that your storage pool is in either site 1 or site 2, then retry the command.

CMMVC8714E The command failed because the number of IO groups specified was different to the number of storage pools specified.

Explanation: When you create HyperSwap volumes, the number of I/O groups that you specify must equal the number of storage pools that you specify.

User response: Retry the command, specifying an equal number of I/O groups and storage pools.

CMMVC8715E The command failed because the storage pools for the volume are not in two independent sites. Volumes can only be created using storage in site 1 and site 2.

Explanation: When you create stretched or HyperSwap volumes with two storage pools, the pools must be located in independent sites, that is, in both site 1 and site 2.

User response: Ensure that you have storage pools in both site 1 and site 2, then retry the command.

CMMVC8716E The create failed because there are no iogroups in site *site_number*.

Explanation: When no I/O group is specified, the system automatically selects an I/O group for the volume in the same site as the MDisk or storage pool that was specified in the command. However, an attempt was made to create a volume in a site where no non-empty I/O groups existed.

User response: Add at least one non-empty I/O group to the site and then retry the command.

CMMVC8717E The create failed because there are no iogroups supporting compression in site *site_number*.

Explanation: When no I/O group is specified, the system selects the I/O group for the volume automatically. No non-empty I/O groups that support compression were found in the site of the specified MDisk or storage pool.

User response: Complete one of the following tasks:

- Specify an existing I/O group that supports compression.
 - Specify an MDisk or storage pool in a site with a non-empty I/O group that supports compression.
 - Reconfigure your system so that the specified MDisk or storage pool is in a site that contains a non-empty I/O group that supports compression.
-

CMMVC8718E The command failed because the managed disk must be in site 1 or 2.

Explanation: Under a stretched or a HyperSwap topology, the managed disk must be in site 1 or site 2.

User response: Ensure that the managed disk is in site 1 or site 2 and retry the command.

CMMVC8719E The command failed because one or more service IP addresses aren't set.

Explanation: The quorum application requires a connection to all nodes in the cluster by using the service IP addresses.

User response:

1. Determine which IP version needs to be used by the quorum application to connect to the cluster: IPv4 (default) or IPv6. To change the default IP version, use the **mkquorumapp -ip_v6** command.
 2. The service IPs of the IP version you are not using might be blank and are not used.
 3. Check that the IPv4 or IPv6 service IPs of all nodes in the cluster are set.
 4. Set any blank IPv4 or IPv6 service IPs. To set a service IP, use the **satask chserviceip** command.
 5. Reissue the command.
-

CMMVC8720E The command failed because the IP address provided is already configured as a DNS server.

Explanation: Each configured DNS server must have a different IP address.

User response: Retry the command with a different IP address.

CMMVC8721E The command failed because the IP address provided is already in use by this system.

Explanation: The IP address that is specified in the command is an IP address on the current system and the system does not provide a DNS service.

User response: Provide the correct IP address of a DNS server.

CMMVC8722E The command has failed because DNS server cannot be removed while DNS clients are configured.

Explanation: The last configured DNS server can be deleted only if no DNS clients are configured.

User response: Stop all DNS clients for the specified server and then retry the command. Because different kinds of DNS clients can trigger this message, specific instructions for stopping the clients cannot be provided here. Refer to the product documentation or contact

your service support representative.

CMMVC8723E The command failed because the system does not support the cloud gateway function.

Explanation: An attempt was made to enable the cloud gateway function when at least one node in the system does not support the function.

User response: Check the hardware compatibility matrix in the documentation. Remove the nodes that do not support the function from the system and try again.

CMMVC8724E The command failed because the candidate does not support cloud snapshots and there is a configured cloud account.

Explanation: An attempt was made to add a node or enclosure that does not support cloud snapshots to a system on which the cloud snapshot function is enabled.

User response: Add a different node, or disable cloud snapshots by using the **rmcloudaccount** command and try again.

CMMVC8725E The command failed because DNS servers are not configured.

Explanation: An attempt was made to configure a cloud account, but the system does not have any configured DNS servers.

User response: Configure at least one DNS server by using the **mkdnsserver** command, and then retry the **mkcloudaccountawss3** command.

CMMVC8726E The command failed because the encryption capability is not currently available on this system.

Explanation: An attempt was made to configure an encrypted cloud account, but the encryption capability is not currently available on this cluster.

User response: Enable encryption on the system by using the **chencryption** command, then retry the **mkcloudaccount** command.

CMMVC8727E The command failed because the cloud gateway service is offline.

Explanation: An attempt was made to test or modify the properties of a cloud account. The system cloud gateway service has reset too often and is being held offline. Cloud accounts cannot be manipulated when they are in this state.

User response: Check the log for an alert event that is related to this error. Mark the event as fixed, and the

system will try to bring the account online. If the attempt fails, a new error message is displayed. Refer to the documentation for the displayed error code for further instructions, then resubmit the command.

CMMVC8728E The command failed because the cloud account is with a different provider.

Explanation: An attempt was made to modify an existing cloud account, but the wrong command was used. For example, if a cloud account is created that uses AWS S3 as the cloud provider, only the **chcloudaccountawss3** command can be used to modify it.

User response: Use the command that is appropriate to the account type to modify an account.

CMMVC8729E The command failed because the maximum number of systems are already using the cloud account.

Explanation: An attempt was made to configure a cloud account with credentials that identify cloud storage that is already being used by the maximum number of systems.

User response: Specify different cloud storage or stop some systems from using the specified cloud storage.

CMMVC8730E The command failed because of a problem accessing metadata objects in the cloud storage.

Explanation: An attempt was made to access cloud storage that was written by a different system, or by the current system at a different time. The current system cannot correctly read the metadata that was written to the cloud storage by the originating system.

User response: Verify with your cloud service provider that the cloud storage is working correctly. If you are trying to access cloud storage that was written by a different system, verify that the system is working correctly. Retry the command.

CMMVC8731E The command failed because the metadata objects were created by a more recent version of code that is incompatible with this system.

Explanation: An attempt was made to access cloud storage that was written by a different system, or by the current system at a different time. The metadata that was created by the other system is not compatible because the other system is running a more recent level of code than the current system.

User response: Update the code level of the current system to match that of the other system, then retry the command.

CMMVC8732E The command failed because the new credentials identify different cloud storage.

Explanation: An attempt was made to change the login details of a cloud account. The new details work, but the cloud storage resources that are accessible by using the new details are different from the storage resources that the system is using.

User response: If you want to use different cloud storage resources, set up a new cloud account object. If you do not want to use different resources, check the details that you are trying to change and try the command again.

CMMVC8733E The command failed because there is at least one cloud account configured.

Explanation: Ordinarily, this error is seen only by service support representatives. An attempt was made to modify the cluster ID alias. At least one cloud account is configured that is using the current cluster ID alias to identify data in the cloud. The cluster ID alias cannot be changed in this situation because data exists in cloud storage that refers to it.

User response: If you specified the correct cluster and must change the cluster ID alias, you must first delete the associated cloud account by using the **rmcloudaccount** command. You can then retry the **chsystem** command.

CMMVC8734E The command failed because the system cannot resolve the hostname of the cloud provider.

Explanation: The system cannot resolve the host name that is associated with the cloud provider by using the configured DNS servers for the system. For some cloud providers, the user entered this host name. For others, such as Amazon S3, the system determines the host name.

User response: If an endpoint URL that contains a host name is associated with the cloud account, verify that the host name is correct. Verify that the system DNS servers are configured correctly and are working correctly. After you correct any problems that you find, retry the command.

CMMVC8735E The command failed because the system cannot contact the cloud provider server on the management network.

Explanation: The system cannot contact the cloud provider over the management Ethernet network.

User response: If an endpoint URL that contains an IP address is associated with the cloud account, verify that the IP address is correct. Determine why the system management IP address cannot contact the cloud

provider; a firewall might be in the way, or the cloud account provider might be experiencing problems. After you correct any problems that you find, retry the command.

CMMVC8736E The command failed because the system cannot establish a connection with the cloud provider software.

Explanation: The system cannot communicate with the cloud provider server software. The cloud provider address on the system might be configured incorrectly or the cloud provider might be having a problem.

User response: If an endpoint URL is associated with the cloud account, verify that the URL is correct. Check that the cloud provider is functioning correctly and is compatible with the system. After you correct any problems that you find, retry the command.

CMMVC8737E The command failed because the system does not have a CA SSL certificate for the cloud provider server.

Explanation: An attempt was made to configure a cloud account that uses SSL. The cloud provider server presented a certificate, but the system cannot verify its authenticity because the system does not have a corresponding certificate authority (CA) certificate.

User response: Contact the cloud provider and obtain the appropriate CA certificate. Or, if SSL is not a requirement, see whether the cloud provider will support a connection where SSL is disabled.

CMMVC8738E The command failed because the SSL certificate has expired.

Explanation: An attempt was made to configure a cloud account that uses SSL, and an SSL certificate was provided, but the current system time is not within the validity period for the certificate.

User response: Verify that the system time is correct. If so, contact the cloud provider to obtain a valid SSL certificate.

CMMVC8739E The command failed because the SSL certificate is not valid.

Explanation: An attempt was made to configure a cloud account that uses SSL. The SSL certificate that was supplied does not work because it is not valid or because it uses features that the system does not support.

User response: Verify that the certificate that you supplied is a valid X509 SSL CA certificate in PEM format. If not, replace the certificate. Verify that the system security setting is compatible with the cloud provider. After you correct any problems that you find, retry the command.

CMMVC8740E The command failed because the account credentials were rejected by the cloud provider.

Explanation: The cloud provider did not accept the login credentials that were configured on the local system.

User response: Verify that the time is set correctly on the system. Verify that the credentials that were entered on the system match valid credentials on the cloud provider. After you correct any problems that you find, retry the command.

CMMVC8741E The command failed because the account credentials do not give permission to access the cloud storage containers.

Explanation: The cloud provider accepted the login credentials that were configured on the local system, but they do not give the system enough permissions to use cloud storage.

User response: Verify that the container prefix that is associated with the cloud account is available and not in use by other cloud clients. Either change the credentials so that they specify a user that has more permissions, or access the cloud provider and give the current user the necessary permissions, as described in the system documentation.

CMMVC8742E The command failed because of an error communicating with the cloud provider.

Explanation: The system encountered an unexpected error when it attempted to communicate with the cloud provider.

User response: Verify that the cloud provider is working normally. Check the system event log for alerts. If you cannot determine the source of the problem, contact your service support representative.

CMMVC8743E The command failed because the region specified is not valid.

Explanation: An attempt was made to configure an Amazon Web Services (AWS) S3 cloud account, but the specified AWS region was not recognized.

User response: Retry the command with a valid region. To find a valid region, check the list of regions that are supported by AWS S3 on the AWS website:

<https://aws.amazon.com/about-aws/global-infrastructure/regional-product-services/>

CMMVC8744E The command failed because the bucket prefix is already in use.

Explanation: An attempt was made to configure a cloud account that uses Amazon Web Services, but the bucket prefix was already in use.

User response: Check the Amazon Web Services documentation on bucket naming. Choose another bucket prefix and retry the command.

CMMVC8745E The command failed because the system is in Gen1 compatibility mode.

Explanation: You cannot configure a cloud account when the system is in Storwize V7000 Gen1 compatibility mode.

User response: Complete the following steps:

1. Ensure that no Storwize V7000 Gen1 canisters are attached to the system.
2. Disable Gen1 compatibility mode by entering the `chsystem -gen1compatibilitymode no` command.
3. Retry the `mkcloudaccountawss3` command.

CMMVC8746E The command failed because an account with an SSL certificate must have an https endpoint.

Explanation: If you provide an SSL certificate when you configure a cloud account, you must provide an endpoint URL that requires SSL. If you provide an endpoint that does not use https, you must not provide an SSL certificate.

User response: Retry the command, either with an https endpoint URL or without the SSL certificate.

CMMVC8748E The command failed because the cloud data is encrypted with the wrong key.

Explanation: The master key associated with the cloud data does not match the master key that was used when the system was created. Cloud snapshot services remain unavailable until this problem is fixed.

User response: Locate the correct master key and make it available to the cluster, either on a USB drive or on a network key server. Then, retry the command.

CMMVC8749E The command failed because the specified volume group is not empty.

Explanation: An attempt was made to delete a volume group that contains volumes. Volume groups cannot be deleted while they contain volumes.

User response: Remove the volumes from the group by using the `chvdisk -novolumegroup` command and try again.

CMMVC8750E The command failed because the specified volume group already contains the maximum number of volumes.

Explanation: The volume cannot be added to the specified volume group after its maximum volume member limit is reached.

User response: Try one of the following solutions:

- Use the specified volume as a stand-alone volume.
- Remove an existing volume from the group by using the **chvdisk -novolumegroup** command, then try again to add the new volume.
- Retry the command and specify a different volume group.

CMMVC8751E A new cloud snapshot could not be created because the volume group does not contain any volumes.

Explanation: An attempt was made to back up an empty volume group.

User response: Add at least one volume to the group and try again. Alternatively, specify a different volume group.

CMMVC8752E A cloud snapshot could not be created because cloud snapshots are not enabled for all the volume group members.

Explanation: An attempt was made to back up a volume group where at least one volume in the group does not have the cloud snapshot function enabled. A new cloud snapshot can be created only when the cloud backup function is enabled for all the volume group members.

User response: Enable the cloud snapshot function for all the volumes in the group and try again. Use the following command to enable cloud snapshots for a volume:

```
chvdisk -backup cloud -enable -account account
volume
```

CMMVC8753E A new cloud snapshot could not be created because a cloud snapshot operation is already in progress.

Explanation: A new cloud snapshot can be started only when the volume **backup_status** is **ready**.

User response: Wait for the existing cloud snapshot operation to complete, or cancel the cloud snapshot. Use the management GUI or the **lsvolumebackupprogress** command to monitor the progress of the cloud snapshot, or the **rmvolumebackupgeneration** command to cancel the cloud snapshot.

CMMVC8754E A cloud snapshot could not be created because a restore operation is already in progress.

Explanation: A new cloud snapshot cannot be started while the volume is being restored.

User response: Make sure that you specified the correct volume. If so, wait for the existing restore operation to complete, or cancel the restore. Use the management GUI or the **lsvolumerestoreprogress** command to monitor the progress of the restore, or the **restorevolume** command to cancel the restore.

CMMVC8755E The command failed because the system does not support the volume group function.

Explanation: An attempt was made to enable the volume group function, but at least one node in the system does not support the function.

User response: Check the hardware compatibility table in the documentation. Remove the nodes that do not support the function from the system and try again.

CMMVC8756E A new cloud snapshot could not be created because the volume is part of a volume group.

Explanation: A new cloud snapshot can be created only if the volume is not part of a volume group.

User response: Remove the volume from the group or use the **backupvolumegroup** command to create a cloud snapshot of the entire group.

CMMVC8757E Cannot change facility because CADF reporting is set.

Explanation: An attempt was made to change the facility code for the syslog server. The facility code is automatically set to 8 when Cloud Auditing Data Federation (CADF) reporting is enabled.

User response: To change the facility code, you must first disable CADF reporting by using the **chsyslogserver -cadf off** command.

CMMVC8758E Cloud snapshots cannot be enabled because there are insufficient free volume copy IDs in the system.

Explanation: The cloud snapshot function requires IDs to create two internal volumes in the same caching I/O group as the volume.

User response: Remove unwanted volumes or volume copies from the system by using the **rmvolume** or **rmvolumecopy** command, and then retry the command.

CMMVC8759E Cloud snapshots cannot be enabled because there are insufficient free volume copy IDs in the caching I/O group of the volume.

Explanation: The cloud snapshot function requires IDs to create two internal volumes in the same caching I/O group as the volume.

User response: Complete one of the following actions:

- Remove unwanted volumes from the caching I/O group.
- Move the volume to a different I/O group.

Then, retry the command.

CMMVC8760E Cloud snapshots cannot be enabled because there is insufficient free capacity in the storage pool.

Explanation: The cloud snapshot function requires added storage capacity in the storage pool of the volume.

User response: Complete one of the following actions:

- Add capacity to the storage pool by using the **addmdisk** command, or use the **mkarray** command if using internal drives.
 - Remove unwanted volumes from the storage pool by using the **rmvdisk** command.
 - Move the volume to a different storage pool by using the **migratevdisk** command.
-

CMMVC8761E Cloud snapshots cannot be enabled because the maximum number of FlashCopy mappings in the system has been reached.

Explanation: The cloud snapshot function configures two internal FlashCopy mappings per volume.

User response: Remove unwanted FlashCopy mappings from the system and retry the command.

CMMVC8762E Cloud snapshots cannot be enabled because there is insufficient memory available for FlashCopy.

Explanation: The restore operation requires an internal FlashCopy mapping to be created, but the system limit was reached.

User response: Remove unwanted FlashCopy mappings from the system by using the **rmfcmap** command and retry the command.

CMMVC8763E Cloud snapshots cannot be enabled because this volume is part of a FlashCopy mapping.

Explanation: The cloud snapshot function cannot be

used with volumes that are the source or target of a FlashCopy mapping.

User response: Make sure that you specified the correct volume. If so, remove all FlashCopy mappings that include this volume and retry the command.

CMMVC8764E Cloud snapshots cannot be enabled because this volume is part of a Metro Mirror or Global Mirror relationship.

Explanation: The cloud snapshot function cannot be used with volumes that are in a remote copy relationship.

User response: Make sure that you specified the correct volume. If so, remove the remote copy relationship that includes this volume and retry the command.

CMMVC8765E Cloud snapshots cannot be enabled because this is a HyperSwap volume.

Explanation: The cloud snapshot function cannot be used with HyperSwap volumes.

User response: Make sure that you specified the correct volume. If so, convert the volume to a basic volume by removing a copy at one site, then retry the command.

CMMVC8766E Cloud snapshots cannot be enabled because the volume is mirrored between two different pools.

Explanation: The cloud snapshot function cannot be used with volumes that have copies in different pools.

User response: Make sure that you specified the correct volume. If so, convert the volume to a basic volume by removing a copy, then retry the command.

CMMVC8767E Cloud snapshots cannot be enabled because the volume is being migrated between pools.

Explanation: An attempt was made to enable the cloud snapshot function while the volume was being migrated between pools.

User response: Wait for the migration operation to finish before you enable cloud snapshots for this volume. Use the management GUI or the **lsmigrate** command to monitor the progress of the migration.

CMMVC8768E Cloud snapshots cannot be enabled because the volume capacity is zero bytes.

Explanation: The cloud snapshot function cannot be used with volumes that have a capacity of zero bytes.

User response: Make sure that you specified the

correct volume. If so, expand the size of the volume and retry the command.

CMMVC8769E Cloud snapshots cannot be disabled because there is a cloud snapshot of the volume in progress.

Explanation: An attempt was made to disable the cloud snapshot function while a cloud snapshot was in progress.

User response: Wait for the cloud snapshot operation to complete or cancel the cloud snapshot. Use the management GUI or the **lsvolumebackupprogress** command to monitor the progress of the cloud snapshot, or the **rmvolumebackupgeneration** command to cancel the cloud snapshot. When the snapshot is no longer in progress, retry the command.

CMMVC8771E The specified parameter lists must be of equal length.

Explanation: The **driveclass** parameter must contain the same number of colon-separated list elements as the **drivecount** parameter.

User response: Validate the command line and retry the command.

CMMVC8772E The command cannot be completed because the I/O group of the drive class has the maximum number of arrays configured.

Explanation: The **mkdistributedarray** command fails because the maximum number of arrays already exists, as configured in the I/O group of the drive class.

Note: Each distributed array occupies 16 slots, starting at an MDisk ID divisible by 16, in the MDisk table. For more information, refer to the documentation for the **lsmdisk** command.

User response: Remove existing arrays in the I/O group before trying the command again.

CMMVC8773E Unable to create the array due to insufficient available drives.

Explanation: The command fails because no drives of the specified drive class in the appropriate drive state are available to create the full array.

Some drives might be available, but not all of the required drives.

Perhaps some drives are not in the candidate state.

Perhaps an enclosure might be offline.

User response: After correcting the drive availability problem for the proper class, use the **lsdriveclass** command or the **lsdriveclass** view in the GUI to verify

the number of available drives. Then try the command again.

CMMVC8774E An array cannot be created with the drive count, stripe width and rebuild areas entered.

Explanation: The **mkdistributedarray** command fails because the values for the drive count, stripe width, and rebuild areas do not work together.

The drive count value must equal or be greater than the combined value of the stripe width and rebuild areas count.

User response: After determining what the proper values should be, try the command again to create the array.

CMMVC8775E Command cannot be initiated because the distributed array does not have an available rebuild space. Replace failed member drives or swap members occupying rebuild spaces before retrying.

Explanation: The **chararraymember** or the **chdrive** command fails because not enough rebuild space is available.

A distributed array has a limited number of spaces where it rebuilds. The command requires the use of a rebuild space, but there are no free spaces to use.

Replacing failed drives creates space in the array to begin a copyback that copies data back from rebuild space. Replacing failed drives thus frees rebuild space for the command to use.

User response: Verify that the array member ID is correct. Replace failed hardware to free rebuild space. If failing a drive, allow the command to degrade redundancy. Then try the command again.

CMMVC8776E A parameter is not supported for this array type.

Explanation: The **chararraymember** or the **chararray** command fails because you used a parameter that is not supported for the array type (distributed or traditional) to attempt a change that is not supported.

Changes to the **rebuildareagoal** and **initnewextents** parameters are supported for distributed arrays only.

Changes to the **sparegoal** and **balanced** parameters are supported on traditional arrays only.

User response: Try the command again using parameters that are appropriate for the type of array.

CMMVC8777E The command failed because the maximum number of encryption keys has been reached.

Explanation: The resources that are needed to create a new key have all been used.

User response: If your object has an individually selectable encryption attribute, consider selecting the **-encrypt no** option for that object. Otherwise, review all encrypted objects in the system. If possible, free up some existing encrypted objects and then retry the command.

CMMVC8778E The command failed because the object is the last active copy of its volume.

Explanation: The volume must contain at least one copy.

User response: If you must delete the copy, delete the entire volume by using the **rmvolume** command.

CMMVC8779E The provided information was not sufficient to identify a specific volume copy.

Explanation: The command uses one or more of the following criteria to delete a volume copy:

- Site number
- Storage pool
- Copy ID

The command did not specify enough information to uniquely identify the volume copy to be deleted.

User response: Retry the command, providing more information.

CMMVC8780E The volume copy was not deleted because it is part of a FlashCopy mapping.

Explanation: You cannot remove a volume copy that is part of a FlashCopy mapping.

User response: First, verify that you specified the correct volume copy. If not, retry the command with a volume copy that is not part of a FlashCopy mapping.

If you specified the correct volume copy, complete one of the following tasks:

- Remove the FlashCopy mapping and then retry the command.
- Specify the **-removefcmaps** parameter to force-delete the volume copy. This parameter stops any dependent FlashCopy maps and must be used with caution.

CMMVC8781E The volume copy was not deleted because it would result in inconsistent data on an image mode mdisk.

Explanation: You cannot delete a volume copy that is out of sync with an image mode MDisk.

User response: First, verify that you specified the correct volume copy. If not, retry the command with a volume copy that does not have an image mode MDisk.

If you specified the correct volume copy, complete one of the following tasks:

- Synchronize the image mode MDisk with the volume copy and then retry the command.
- Specify the **-discardimage** parameter to force-delete the volume copy. A forced deletion can result in data loss on the image mode copy.

CMMVC8782E The volume copy was not deleted because the copy has an image mode mdisk and a change volume is providing the consistent image.

Explanation: A volume copy cannot be deleted when a change volume is in use.

User response: First, verify that you specified the correct volume copy. If not, retry the command with a volume copy that does not have an image mode MDisk.

If you specified the correct volume copy, complete one of the following tasks:

- Wait for the use of the change volume to complete and then retry the command.
- Specify the **-discardimage** parameter to force-delete the volume copy. A forced deletion can result in data loss on the image mode copy.

CMMVC8783E The volume copy was not deleted because the volume is part of a consistency group.

Explanation: A volume copy cannot be deleted when it is part of an active-active relationship.

User response: First, verify that you specified the correct volume copy. If not, retry the command with a volume copy that is not part of a consistency group.

If you specified the correct volume copy, remove the active-active relationship for the volume from the consistency group, then retry the command.

CMMVC8784E The command failed because an additional feature activation is required for the candidate node

Explanation: The candidate node requires an additional license before the action can be completed.

User response: Activate the feature for the candidate node and then retry the command.

CMMVC8785E The command failed because the node does not support encryption, and encryption is required in the IO group.

Explanation: The existing I/O node is capable of encryption. This error occurs when adding an I/O group partner node that does not support encryption.

User response: Acquire an encryption capable node to add to the I/O group.

CMMVC8786E The node cannot be added because it does not support encryption, and encryption is in use for some SAN MDisks.

Explanation: Encryption is enabled. An attempt was made to add a node that is not encryption-capable while storage pools were being encrypted that include SAN MDisks that are not self-encrypting.

User response: Acquire an encryption capable node, or remove all non-self-encrypting SAN MDisks from all encrypting storage pools.

CMMVC8787E The volume was not deleted because it is mapped to a host.

Explanation: You cannot delete a volume that is mapped to a host.

User response: If you specified the wrong volume, retry the command. If you are sure that you want to remove the specified volume, remove the host mapping and retry the command. Alternatively, you can include the **-removehostmappings** parameter with the **rmvolume** command to force the removal of the host mappings.

CAUTION:

When you use a "force" parameter such as **-removehostmappings**, you risk loss of data.

CMMVC8788E The volume was not deleted because it is part of Metro Mirror or Global Mirror relationship.

Explanation: You cannot delete a volume that is part of a remote copy relationship.

User response: If you specified the wrong volume, retry the command. If you are sure that you want to remove the specified volume, remove the remote copy relationship and retry the command. Alternatively, you can include the **-removercrelationships** parameter with the **rmvolume** command to force the removal of the remote copy relationship.

CAUTION:

When you use a "force" parameter such as **-removercrelationships**, you risk loss of data.

CMMVC8789E The volume was not deleted because it is part of a FlashCopy mapping.

Explanation: You cannot delete a volume that is part of a FlashCopy mapping.

User response: If you specified the wrong volume, retry the command. If you are sure that you want to remove the specified volume, remove the FlashCopy mapping and retry the command. Alternatively, you can include the **-removefcmaps** parameter with the **rmvolume** command to force the removal of the FlashCopy mappings.

CAUTION:

When you use a "force" parameter such as **-removefcmaps**, you risk loss of data.

CMMVC8790E The volume was not deleted because it would result in inconsistent data on an image mode mdisk.

Explanation: You cannot delete an image mode volume that has inconsistent data.

User response: If you specified the wrong volume, retry the command. If you are sure that you want to remove the specified volume, wait until the data on the image mode volume is consistent before you retry the command. Alternatively, you can include the **-discardimage** parameter with the **rmvolume** command to force the removal of the volume.

CAUTION:

When you use a "force" parameter such as **-discardimage**, you risk loss of data.

CMMVC8791E There is already an outstanding certificate request. Use the **-force** flag to discard this request and generate a new certificate request.

Explanation: You issued 'chsystemcert -mkrequest' twice without successfully installing the certificate from the first request that uses 'chsystemcert -install'.

User response: Do one of the following:

- Get the outstanding request signed and install it using 'chsystemcert -install'.
- Abandon the outstanding request and start a new one using 'chsystemcert -mkrequest ... -force'.

Note: If a new request is started, any certificate that is created from a previous request can no longer be used.

CMMVC8792E Cannot install certificate because there is no outstanding certificate request.

Explanation: You are attempting to install a certificate file using 'chsystemcert -install' without first creating a certificate request. Only a signed certificate that is produced from a certificate request that is created by the system can be installed. Once the certificate has been installed, it cannot be installed again.

User response: Only a signed certificate request can be installed. Use the 'chsystemcert -mkrequest' command to create a certificate request and then get this request signed by a certificate authority (CA). The resulting certificate can then be installed.

CMMVC8793E Cannot install certificate because it has the wrong key.

Explanation: You are attempting to install a certificate file by using 'chsystemcert -install', which is not derived from the outstanding certificate request. Either:

- You tried to install the wrong certificate
- 'mkrequest -force' was used between the certificate request being generated and the installation of the signed certificate.

Note: Get the latest request signed and install this certificate.

User response: Find the correct certificate to install or use the 'chsystemcert -mkrequest -force' command to start a new certificate request.

CMMVC8794E Invalid certificate file.

Explanation: The certificate file that you are attempting to install cannot be parsed. The certificate must be provided in base64 encoded PEM format.

User response: Check that the certificate to be installed is in the correct format and then copy the file to the system. Attempt to install it again.

CMMVC8795E This command is no longer supported. Use the CLI command 'chsystemcert' to manage the SSL certificate for the system.

Explanation: The CLI command 'chsystem -regensslcert' is no longer supported. A new CLI command 'chsystemcert' that provides more options for managing the SSL certificate for the system, replaced it.

User response: Use the 'chsystemcert' command instead.

CMMVC8796E Cannot make unencrypted MDisk Group as parent MDisk Group [%1] has an encryption key.

Explanation: You are not allowed to make an unencrypted child MDisk Group when the parent MDisk Group has an encryption key.

User response: The use case is to make an unencrypted child MDisk Group on an encrypted system. You will need to make a new MDisk Group without an encryption key instead.

CMMVC8797E Unable to change mdisk property while being part of an MDisk Group with an encryption key.

Explanation: The MDisk Group and/or one of its child pools has an encryption key that it is using for the MDisk. This can not be changed while the MDisk contains customer data.

User response: You need to remove the MDisk from the MDisk Group, apply the change in property, and then add back to the MDisk Group.

CMMVC8798E Cannot add unencrypted MDisk to encrypted MDisk Group

Explanation: The user is not allowed to add an unencrypted MDisk to an encrypted storage pool.

User response: Only MDisk in encrypting I/O groups can be added to an encrypted storage pool.

CMMVC8799E It is not possible to create an encrypted image disk.

Explanation: You are trying to create an image mode disk that would be encrypted.

User response: Image mode disks can be created in an unencrypted storage pool only.

CMMVC8800E Cannot migrate as pool *storage_pool* is encrypted.

Explanation: Encrypted data cannot be moved unless the source and destination have the same encryption key.

User response: Create or add a volume copy and then (auto) delete the source once copied.

CMMVC8801E Insufficient drive count specified.

Explanation: The `lspotentialarraysize` command fails when its `drivecount` parameter is less than its `stripewidth` parameter.

The number of drives in the array must be an integer from 3 - 128. The stripe width varies per RAID type:

RAID Stripe width

R1	2-16
R5	3-16
R6	5-16
R10	Even integers from 2-16

User response: Retry the command with a drive count that is equal to or greater than the stripe width.

CMMVC8802E The command failed because the parent pool contains an unencrypted external MDisk.

Explanation: An attempt was made to create an encrypted child pool where the parent pool has an unencrypted external MDisk and at least one I/O group contains hardware that is not reported as encryption capable. It is not possible for non-encryption-capable nodes to perform encrypted operations.

User response: Either remove the non-encryption-capable I/O group members from the cluster or create the pool as not encrypted.

CMMVC8803E Stripe width incorrect for RAID type.

Explanation: The `lsptentailarraysize` command fails because the stripe width is not good. Valid values are:

RAID	Stripe width
R0	1-8
R1	2
R5	3-16
R6	5-16
R10	Even integers from 2-16

User response: Retry the command with a valid stripe. Use the `lsarrayrecommendation` command to recommend an optimal value.

CMMVC8804E The number of rebuild areas specified must be greater than or equal to the rebuild area goal parameter.

Explanation: The command fails because it creates a distributed array with a goal of more rebuild areas than are available in this array.

User response: If you use the `rebuildareas` parameter of the `mkdistributedarray` command, ensure that its value is greater than or equal to the `rebuildareagoal` parameter value.

If you do not use the `rebuildareas` parameter, but do use the `rebuildareagoal` parameter, the goal value is greater than the default areas value of 1 for RAID5 or 2 for RAID6. Either reduce the goal or specify a greater areas value than the default.

CMMVC8805E An array of this RAID level cannot be created using this drive capacity.

Explanation: The command fails because it cannot create a distributed array with drives that have a capacity exceeding the limit for the RAID level. For example, you cannot create a RAID5 array if the drive capacity is greater than or equal to 8 TB.

User response: Change the RAID level or use drives with a capacity less than the limit for the RAID level, such as 8 TB for RAID5.

CMMVC8806E The specified drive count exceeds available drives in the drive class.

Explanation: The `lsarrayrecommendation` command fails because the colon-delimited list of drive counts has a drive count that is greater than the number of available drives in the respective drive class.

User response: Try the command again with each drive count equal to or less than the number of available drives in its respective drive class.

CMMVC8807E A duplicate drive class ID was entered in the list.

Explanation: The `lsarrayrecommendation` command fails when the colon-delimited list of drive classes has duplicate IDs.

User response: Try the command again after editing the list of drive classes.

CMMVC8808E Cannot create the encrypted child pool because the parent pool contains an array that is provided by a non-encryption capable IO group.

Explanation: You attempted to create an encrypted child pool where the parent pool contains an array that is provided by an enclosure that does not support encryption.

User response: Either remove the array from the parent group, create the child pool with a parent pool that doesn't contain such arrays or create the child pool without encryption.

CMMVC8809E Cannot create the encrypted child pool because the parent pool contains an unencrypted DRAID array.

Explanation: You attempted to create an encrypted child pool where the parent pool contains an unencrypted DRAID array.

User response: Remove the array from the parent group or create the child pool without encryption.

CMMVC8810E The volume was not deleted because it is configured as a change volume for a Remote Copy relationship. Either disassociate the volume from the relationship or use **-removefcmaps** to force delete the volume, which might result in the loss of data.

Explanation: An attempt was made to delete a volume that is configured as the master or auxiliary change volume in a remote copy relationship.

User response: Complete one of the following actions:

- If you entered the wrong volume, correct the command and retry it.
- Remove the volume from the remote copy relationship and retry the command.
- Retry the command with the **-removefcmaps** parameter to force the deletion of the volume.

CAUTION:

A forced deletion can result in data loss.

CMMVC8811E The command is not supported for this array type.

Explanation: Any of the following actions can cause this error message to be displayed:

- Trying to change a RAID array to be encrypted by using the **charray -encrypt** command.
- Trying to change the balanced or spare goal parameters for a DRAID array.
- Trying to change the rebuild areas goal or init new extents parameters for a TRAID array.
- Trying to change the balance parameter for a DRAID member.
- Trying to change a DRAID array to use latency mode.
- Trying to create a DRAID array in latency mode.
- Running the **lspotentialarraysize** command for RAID0
- Running the **lspotentialarraysize** command for RAID1 if rebuild areas set or for RAID10 if rebuild areas set.

User response: Specify an array of the correct type and retry the command.

CMMVC8812E The command failed because the existing volume may have only one mirrored copy.

Explanation: In a HyperSwap topology, the existing volume can have only a single mirrored copy.

User response: Specify a different volume or remove an existing mirrored copy, and then retry the command.

CMMVC8813E The command failed because the existing volume has mirrored copies in different sites.

Explanation: In a HyperSwap topology, all of the mirrored copies for the existing volume must have a matching site of 1 or 2 (not 3 or none). This configuration ensures that the existing volume itself has a consistent and well-defined site of 1 or 2.

User response: Ensure that all mirrored copies of the existing volume have a matching site of 1 or 2 and then retry the command.

CMMVC8814E The command failed because the existing volume has mirrored copies in storage pools whose site is not 1 or 2.

Explanation: In a HyperSwap topology, all of the mirrored copies for the existing volume must have a matching site of 1 or 2 (not 3 or none). This configuration ensures that the existing volume itself has a consistent and well-defined site of 1 or 2.

User response: Ensure that all mirrored copies of the existing volume have a matching site of 1 or 2 and then retry the command.

CMMVC8815E The command failed because the IO group is not in the same site as the storage pool.

Explanation: The **-iogrp** parameter applies only when you create a HyperSwap volume copy and requires that the system be configured with a HyperSwap topology. The caching I/O group must be in the same site as the storage pool that you specify for the volume copy that you are creating.

User response: Correct the command and retry.

CMMVC8816E The command failed because the specified FlashCopy mapping is controlled by FlashCopy.

Explanation: This command is not valid for an internally controlled FlashCopy map.

User response: Specify a FlashCopy map that is not internally owned and retry the command.

CMMVC8818E The command failed because the volume already has a copy in the specified site. Use the **addvdiskcopy** command to add a second copy in the same site.

Explanation: An attempt was made to add a volume copy to a site by using the **addvolume copy** command when the volume already has a copy in the specified site.

User response: Use the **addvdiskcopy** command to

create an additional copy in the same site.

CMMVC8819E The volume copy could not be created because the existing source copy is not compressed.

Explanation: When creating new volume copies, use the `ignoresyncerrors` parameter only when the existing source copy is a compressed copy.

User response: Use the `ignoresyncerrors` parameter for compressed volume copies.

Note: The `ignoresyncerrors` parameter can also help create valid copies of compressed volume copies that were corrupted due to a bug in the compression engine.

CMMVC8870E The volume was not created because a capacity of zero bytes is not allowed for Hyperswap volumes.

Explanation: A HyperSwap volume can only be created with a virtual capacity greater than zero bytes.

User response: Either specify a non-zero size for the volume or create a different type of volume.

CMMVC8879E I/O group specified is not present.

Explanation: If you must perform discovery by using a specified I/O group, then you must specify a valid I/O group number. An invalid or non-existent I/O group was specified.

User response: Obtain a valid I/O group number by using the management GUI or the `lsiogroup` command.

CMMVC8880E Some of the nodes in the cluster/specified IO group are not in online state.

Explanation: One or more of the nodes that are part of the specified system or I/O group are not in an active state. You must initiate discovery from all of the nodes in the system or in the specified I/O group.

User response: Check the node status by using the management GUI or the `lsnode` or `lsnodecanister` command, and ensure that all nodes are online before you restart discovery.

CMMVC8881E Discovery cannot be run as the source port ID is invalid for one or more nodes in the system or specified iogroup.

Explanation: Different models or generations of nodes might have different numbers of Ethernet ports. If the system is created with multiple nodes that have different supported node models, the source port ID

that you specify (either by using the management GUI or by using the `detectiscsistorageportcandidate` command) cannot be greater than the highest port ID on any node that is part of the system or specified I/O group.

User response: Check the number of Ethernet node ports present on each node by using the management GUI or the `lsportip` command. Run discovery again, and this time specify a port ID within the available range.

CMMVC8882E Some of the specified source Ethernet ports are not configured for iSCSI use.

Explanation: The Ethernet ports on all nodes in the system or specified I/O group must be in an online state before a backend controller discovery can be initiated through the port. The port can be in one of the following states:

unconfigured

No iSCSI address is assigned to the port.

configured

An IP address is assigned to the port.

User response: Use the management GUI or the `lsportip` command to identify the Ethernet port on the unconfigured node whose source port ID you specified for discovery. Use the management GUI or the `cfgportip` command to assign an IP address of the required type (IPv4 or IPv6). Ensure that all required node Ethernet ports are in an online state before you restart discovery.

CMMVC8883E Link state is inactive for one or more source Ethernet ports in the system or specified IO group.

Explanation: The link must be active for all node Ethernet ports that are specified with a source port ID. If any of the links are down, you cannot initiate discovery.

User response: Use the management GUI or the `lsportip` command to identify the node Ethernet port with source port ID whose `link_state` is inactive. Verify that the Ethernet cable is connected to the port for all affected nodes, then restart discovery.

CMMVC8884E IP type mismatch between IP of source ports and target IP specified.

Explanation: To establish an iSCSI session, the initiator and target connection endpoints must both have either IPv4 or IPv6 addresses. If you are performing cluster-wide discovery and the specified target IP type is IPv6, then an IPv6 address must be assigned to all of the initiator source ports with the specified port ID. If you are performing I/O-group-specific discovery and the specified target IP type is IPv6, then an IPv6 address must be assigned to all of the initiator source

ports in the I/O group with the specified port ID.

User response: Use the management GUI or the **cfgportip** command to assign IP addresses with an IP type that matches the target IP. Alternatively, you can configure an IP address that matches the IP type of the source ports on the target controller port.

CMMVC8885E New controller cannot be discovered as limit of managed iSCSI controllers has been reached.

Explanation: The **detectiscsistorageportcandidate** command found a new controller with no existing established connections after the maximum of 64 iSCSI controllers was reached.

User response: Use the management GUI or the **lscontroller** command to view the list of managed controllers. Identify any controllers that you no longer need to manage. Remove the mdiskgroups that are composed of mdisks that were exported from the identified controller. Use the **rmiscsistorageport** command to remove the unused iSCSI port. After the **lscontroller** command shows fewer than the maximum of 64 controllers, retry the command to discover the new controller. Also, consider how you might consolidate iSCSI storage onto fewer controllers.

CMMVC8886E The targetip/ targetip6 parameters do not adhere to the standard IPv4/ IPv6 formats.

Explanation: An IPv4 or IPv6 address was used in a **targetip** or **targetip6** argument that does not adhere to standard IP address format.

User response: Reenter the command, this time using the correct format in the **targetip** or **targetip6** argument. See the following table for accepted formats.

Table 118. Accepted IP address formats

IP type	IP address list format
IPv4	1.2.3.4
Full IPv6	1234:1234:abcd:0123:0000:0000:7689:6576
Full IPv6, leading zeros suppressed	1234:1234:abcd:123:0:0:7689:6576
IPv6 with zero compression	1234:1234:abcd:123::7689:6576

CMMVC8887E Authentication with target controller using the credentials specified did not succeed

Explanation: The **username** and **chapsecret** arguments that the iSCSI initiator presented do not match the credentials expected by the iSCSI target controller.

User response: Make sure that you specify the correct credentials in the **username** and **chapsecret** arguments.

CMMVC8888E Discovery could not start as a connection could not be established to the backend controller.

Explanation: A problem in the IP network or the iSCSI target controller prevented the login or discovery request from reaching the target controller, or prevented the target controller from responding to the discovery request.

User response: Verify the following conditions:

- The IP network is properly configured
- The target ports are configured
- The target IPs are reachable

CMMVC8889E Discovery could not be completed due to problems at target or initiator.

Explanation: A problem occurred that prevented a successful login or prevented discovery output from the target controller. Possible problems include protocol errors; a transient target, initiator, or network problem; or other errors returned by the target in response to the discovery request.

User response: Correct any configuration problems at the initiator or target. If the problem persists, contact IBM Support.

CMMVC8890E Sessions cannot be established as there are no discovered controllers for the specified id.

Explanation: A row ID was not specified that indicates the target iSCSI name.

User response: Use the management GUI or the **lscsistorageportcandidate** command to list the discovered targets and select a valid row ID as an input to the **addiscsistorageport** command.

CMMVC8891E The specified target IP must not refer to the IP of a system management port

Explanation: An IP conflict was detected between IP addresses of source ports and targets. This error prevents discovery requests being sent from one source port to the same or different configured source ports on any node in the cluster.

User response: Use the management GUI or the **lsportip** command to locate any source port IP addresses that are the same as IP addresses configured on the target. Resolve any conflicts and retry discovery.

CMMVC8892E The specified target IP must not refer to an iSCSI port on the system

Explanation: An IP conflict was detected between assigned clustered system IP addresses and target IP addresses. This error prevents discovery requests being sent from one source port to a system management

port. The clustered system IP addresses are assigned during cluster creation.

User response: Resolve conflicting IP addresses and retry the command. If necessary, you can use the management GUI or the **lsystemip** command to locate conflicting IP addresses, and use the **chsystemip** command to change system IP addresses.

CMMVC8893E Sessions cannot be established as discovery has not been done through the iogroup specified.

Explanation: Discovery must be done through an I/O group before you can use the **addiscsistorageport** command to establish sessions through that I/O group. You can run the **addiscsistorageport** command with a specific I/O group ID if you previously ran cluster-wide discovery.

User response: Use the management GUI or the **lsiscsistorageportcandidate** command to view the I/O groups through which discovery was done. Use this information to correct the command, and then resubmit it.

CMMVC8894E Limit of max IQN+IP tuples per external iSCSI storage system has been reached.

Explanation: A maximum of 128 unique IQN + IP combinations (tuples) are allowed per controller from the initiating cluster.

User response: Use the management GUI or the **lsiscsistorageport** command to view the list of established sessions. Remove sessions that are no longer required by using the management GUI or the **rmiscsistorageport** command. Retry the **addiscsistorageport** command.

CMMVC8895E Sessions cannot be established as maximum session count per node has been reached.

Explanation: The maximum of 256 initiator sessions per initiator node were established from one or more target controllers. Each invocation of the **addiscsistorageport** command creates an additional session through the nodes in the specified I/O group or cluster.

User response: Use the management GUI or the **lsiscsistorageport** command to check the total number of sessions per node. Use the management GUI or the **rmiscsistorageport** command to remove any sessions that are no longer in use. When the number of sessions is below the maximum, try again to establish the session.

CMMVC8896E One source port cannot have more than one session to the same target IQN through different target ports.

Explanation: One-to-many connectivity from source ports to backend controller target ports can create bottlenecks in the I/O path.

User response: Use the management GUI or the **lsiscsistorageport** command to view the sessions that were established between initiator ports and the target IQN+IP. Create a session to a target IQN that does not already have a connection to a source port.

CMMVC8897E One node cannot have more than one session with same target IQN+IP through different source ports.

Explanation: Many-to-one connectivity from source ports to backend controller target ports can create bottlenecks in the I/O path.

User response: Use the management GUI or the **lsiscsistorageport** command to view the sessions that were established between initiator ports to target IQN+IP. Use a source port for the session that is not already connected to a target IQN.

CMMVC8898E Sessions cannot be torn down as the session does not exist.

Explanation: A valid existing session identifier was not specified for the tear-down.

User response: Use the management GUI or the **lsiscsistorageport** command to list the sessions that are candidates for removal.

CMMVC8901E Enter at least one throttle parameter.

Explanation: No throttle parameters (such as IOPS or bandwidth) were specified.

User response: Specify one or more throttle parameters.

CMMVC8902E Volume is invalid or does not exist.

Explanation: The ID or name of the volume that is to be throttled was not valid or was not found.

User response: Use a valid ID or name for the volume when you retry the command.

CMMVC8903E Throttle is already associated with this volume.

Explanation: The volume ID or name that you specified in a command line parameter is already associated with a throttle value.

User response: Either specify a different ID or name for the volume, or use the **chthrottle** command to

change an existing throttle value for the current volume.

CMMVC8904E Throttle parameter missing or invalid.

Explanation: You entered an invalid value for one or more of the parameters, or one or more required parameters are missing, or both.

User response: Verify that all required parameters are specified and that all values are valid, and then retry the command.

CMMVC8905E Offload IO throttle already exists.

Explanation: A throttle exists for offloaded I/O. Only one offloaded I/O throttle can be created for each cluster.

User response: To change the offloaded I/O throttle parameters, use the **chthrottle** command.

CMMVC8906E Throttle object does not exist.

Explanation: The value that you specified for the **throttle_name** parameter was not found.

User response: Retry the command, this time specifying a valid throttle name.

CMMVC8910E The specified IP is a reserved or special IP and cannot be used.

Explanation: The following special or reserved IP addresses cannot be assigned to target controllers:

- Broadcast addresses

Note: Because no subnet mask is specified, the system does not detect the incorrect use of subnet broadcast addresses.

- Multicast addresses
- Unspecified addresses ("0.0.0.0" for IPv4, "::" for IPv6)
- Loopback addresses
- Link-local addresses

Use APIPA addresses with caution because the system cannot verify them.

User response: Configure a supported IP address on the target controller and retry discovery.

CMMVC8911E Limit of IQNs per target controller has been reached.

Explanation: A maximum of 64 iSCSI Qualified Names (IQNs) per target iSCSI controller is permitted.

User response: Use the management GUI or the **lsiscsistorageport** command to view the sessions for a controller. Remove sessions to IQNs that are no

longer in use by using the management GUI or the **rmiscsistorageport** command. Rerun the **addiscsistorageport** command.

CMMVC8912E The specified Ethernet port is not configured for connectivity with external storage systems.

Explanation: The Ethernet port must be configured to allow connectivity to a backend storage controller before you can begin discovery or establish a session through the port.

User response: Use the management GUI or the **cfgportip** command to enable storage system connectivity. Use the **storage** flag for IPv4 addresses and the **storage_6** flag for IPv6 addresses.

CMMVC8913E The port IP address cannot be unconfigured as there are active sessions established to backend controller(s) through the port.

Explanation: An attempt was made to remove an IP address from a port that was used to establish backend controller sessions. Removing an IP address through which sessions are established puts mdisks into a degraded state and is not permitted.

User response: Use the management GUI or the **rmiscsistorageport** command to remove sessions through the source port before you remove an IP address that is used for backend controller connectivity. Alternatively, if the intent is only to remove the IP and then configure a new IP, you can achieve the same goal by using the **cfgportip** command with a **-force** flag.

CMMVC8914E Storage flag cannot be reset as the IP is already in use for backend controller(s) connectivity.

Explanation: If you set the **storage** or **storage_6** flag to yes, the IPv4 or IPv6 address for the port can be used to establish connections to the backend iSCSI controller. After you use the **addiscsistorageport** command to establish connectivity from the source port to one or more backend controllers, then you cannot set the **storage** or **storage_6** flag to no unless all the sessions to backend controllers are removed.

User response: Verify the information, make a correction, and resubmit the command.

CMMVC8915E Changing the IP address attributes may cause mdisks to go into a degraded state for some time.

Explanation: If a source IP address is in use for backend controller connectivity, changing the IP address, mask, or gateway causes the existing sessions to be removed and new sessions to be established.

During this phase, the mdisks that are visible though the source port that is being reconfigured go into a degraded state until the new sessions are established.

User response: If you did not intend to reconfigure, reenter the command with the correct parameters. If you do want to reconfigure, an administrator can use the management GUI or the **cfgportip** command with the **-force** flag to go ahead with the reconfiguration.

CMMVC8919E The command failed because the requested transition is invalid.

Explanation: An invalid transition was attempted, such as from disabled to enabled.

User response: Use multiple steps to make the required transition. For example, move from disabled to transitional to enabled. Refer to the help for specific configuration instructions.

CMMVC8920E The command failed because one or more host ports are currently failed over.

Explanation: A state transition was attempted while a failover was in progress. In practice, this error is rarely seen, due to the requirement that both nodes be online. A small timing window for the error can occur after a node has just unpended and before a controlled node shutdown.

User response: Wait until the **lstartportfc** command reports that all ports in this I/O group are online on their owner nodes, and then retry.

CMMVC8922E The command failed because the volume copies would not be in two independent sites.

Explanation: In a stretched system topology, you can add a volume copy only if the result is a volume copy in each of site 1 and site 2.

User response: Ensure the existing volume copy has a well defined **site** argument (site 1 or site 2), and retry by adding the new volume copy to the other site.

CMMVC8923E The command failed because a volume in a Metro Mirror or Global Mirror relationship cannot have copies in two sites using HyperSwap.

Explanation: In a HyperSwap system topology, you cannot add a volume copy to a different site if the existing volume is in a Metro Mirror or Global Mirror relationship.

User response: If you specified the wrong volume, correct the command and retry. If you do need to add the volume copy to a different site, first remove the Remote Copy relationship by using the **rmrcrelationship** command, then retry the command.

CMMVC8924E The command failed because the volume already has a copy in the specified site.

Explanation: In a stretched system topology, a volume copy can be added only if the result is a volume copy in each of site 1 and site 2.

User response: Retry by specifying a storage pool with a different site for the new volume copy.

CMMVC8925E The command failed because the cache state for an image mode volume or copy is corrupt.

Explanation: The data on the image mode mdisk might be different from the data on the image mode volume.

User response: Use the **recovervdisk** command to acknowledge the volume data loss, then retry the command. Alternatively, specify the **-discardimage** parameter to force the removal of the image mode copy.

CMMVC8926E The volume copy was not deleted because a change volume is providing the consistent image for a copy on the other site.

Explanation: An attempt was made to remove an image mode copy of a HyperSwap volume while a change volume was still in use.

User response: Wait until the change volume is no longer in use, then retry the command. Alternatively, specify the **-discardimage** parameter to force the removal of the volume copy, which might result in data loss.

CMMVC8933E Limit of host clusters has been reached.

Explanation: An attempt was made to create a host cluster when the maximum number of host clusters had already been created. The maximum number of host clusters is 512.

User response: Remove one or more existing host clusters before you add a new one.

CMMVC8934E Host already belongs to a host cluster.

Explanation: An attempt was made to add a host to a host cluster when the host already belongs to another host cluster.

User response:

- If you specified the wrong host, retry the command and specify a different host.

- Otherwise, remove the host from the host cluster to which it currently belongs, then retry the command.

CMMVC8935E Host cluster does not exist.

Explanation: A host cluster was specified that does not exist.

User response: Retry the command with a different host cluster.

**CMMVC8936E Host cluster has shared mappings.
Use -keepmappings or
-removemappings flag.**

Explanation: An attempt was made to delete a host cluster that has one or more shared mappings without setting a flag to specify how those mappings are to be treated.

User response: Retry the command and specify one of the following flags:

-removemappings

To delete the shared mappings along with the host cluster.

-keepmappings

To keep the shared mappings as private mappings for the host.

CMMVC8937E Limit of hosts in the host cluster has been reached.

Explanation: An attempt was made to add a host to a host cluster that is full. A host cluster can include a maximum of 128 hosts.

User response:

- Add the host to a different host cluster.
- Alternatively, remove one or more hosts from the host cluster before you add another.

CMMVC8938E Host is not part of the host cluster.

Explanation: An attempt was made to carry out one of the following actions, specifying a host that did not belong to the specified host cluster:

- Remove the host
- Convert a shared mapping to a private mapping

User response: Retry the command and specify a different host, a different host cluster, or both.

CMMVC8939E Some source Ethernet ports are temporarily unavailable due to failover.

Explanation: In the case of a node-pend event, the iSCSI IP addresses assigned to node Ethernet ports fail over to the partner node. While the IP addresses are in a failover state, discovery cannot be initiated from the node those addresses are assigned to, even though the

node ports show as configured. Discovery can be initiated when the IP addresses fail back. The IP addresses fail back immediately if the node unpends 5 minutes or more after the pend event. If the node unpends less than 5 minutes after the pend event, the IP addresses fail back only after pend time + 5 minutes.

User response: Use the management GUI or the **lsportip** command to verify that the IP addresses failed back to the owning node before retrying discovery.

CMMVC8940E Discovery has not been initiated from nodes of specified site.

Explanation: The **addiscsistorageport** command was not run with the same site ID argument that was used to run discovery.

User response: Run the management GUI or the **addiscsistorageport** command with the same site ID argument that was used to run discovery with the **detectiscsistorageportcandidate** command.

CMMVC8941E Sessions cannot be added to the target controller due to conflict with existing connectivity or discovery output.

Explanation: One of the following situations occurred:

- Discovery was run with one value for the **iogroup** parameter and the **addiscsistorageport** command was run with a different value for the **iogroup** parameter.
- Discovery was run with a value for the **iogroup** parameter and the **addiscsistorageport** command was run cluster-wide (**iogroup** not specified).
- Discovery was run cluster-wide and the **addiscsistorageport** command was run with a value for the **iogroup** parameter.
- Sessions were already established through one I/O group, and the **addiscsistorageport** command was run to add connectivity through another I/O group or cluster-wide.
- Sessions were already established cluster-wide, and the **addiscsistorageport** command was run to add connectivity through an I/O group.
- Sessions were already established through a site, and the **addiscsistorageport** command was run to add connectivity through an I/O group or cluster-wide.
- Sessions were already established through an I/O group or were established cluster-wide, and the **addiscsistorageport** command was run to add connectivity through a site. (You cannot mix I/O groups of different topologies in the same system.)

User response: Run the management GUI or the **addiscsistorageport** command with the same site ID parameter that was used to run discovery with the **detectiscsistorageportcandidate** command.

CMMVC8942E The volume copy was not deleted because the volume is part of a FlashCopy mapping. All FlashCopy mappings must be removed from the volume before this copy can be deleted.

Explanation: An attempt was made to remove a volume copy when the volume had FlashCopy mappings.

User response: Remove all FlashCopy mappings from the volume, then try again to delete the copy.

CMMVC8943E The command failed because the consistent image on an image mode copy is being provided by a FlashCopy map.

Explanation: An attempt was made to remove an image mode volume or volume copy that is the target of an active (copying) FlashCopy mapping. While the FlashCopy mapping is copying, the data on the image mode mdisk might be different from the data on the image mode volume.

User response: Wait for the FlashCopy operation to complete, then retry the command. Alternatively, specify the **-discardimage** parameter to force the delete operation, which might result in data loss.

CMMVC8944E The command failed because the volume already has copies in two sites using HyperSwap. Use the **addvdiskcopy** command to add a second copy in the same site.

Explanation: In HyperSwap system topology, the **addvolumecopy** command cannot be used to add a mirrored copy in the same site as an existing volume copy.

User response: Use the **addvdiskcopy** command to add a second copy in the same site.

CMMVC8945E The command failed because it would allow access to an image mode copy in a Metro Mirror or Global Mirror relationship that does not contain consistent data.

Explanation: An attempt was made to remove an image mode volume or volume copy that is a secondary in a Metro Mirror or Global Mirror relationship and the data on the volume is not consistent.

User response: Wait for the relationship to synchronize the volume from the primary volume, then retry the command. Alternatively, specify the **-discardimage** parameter to force the delete operation.

CMMVC8946E The command failed because an image mode copy is not synchronized with the host-accessible copy.

Explanation: An attempt was made to remove an image mode volume or volume copy, and the data on an image mode copy is not synchronized with the host-accessible copy.

User response: Wait for the volume copy to resynchronize, then retry the command. Alternatively, specify the **-discardimage** parameter to force the delete operation.

CMMVC8947E The target controller is not in the storage layer.

Explanation: The target Storwize controller is configured in the replication layer.

User response: Consider whether the target controller is configured correctly. If appropriate, use the management GUI or the **chsystem -layer storage** command on the target controller to configure the controller in the storage layer. If the target controller is in the correct layer, check that the IP address that you specified points to the correct target. Correct the layer for the target controller or specify the correct IP address and rerun discovery by using the **detectiscsistorageportcandidate** command.

CMMVC8948E Storwize V7000 Gen1 compatibility mode cannot be disabled because the system contains a Storwize V7000 Gen1 enclosure.

Explanation: An attempt was made to disable Storwize V7000 Gen1 compatibility mode. This mode must be enabled on any system that contains Storwize V7000 Gen1 enclosures alongside Storwize V7000 Gen2 enclosures. The system contains such a combination, so the request to disable the mode was refused.

User response: If you must disable Storwize V7000 Gen1 compatibility mode, remove the Storwize V7000 Gen1 enclosures from the system and run the command again.

CMMVC8949E Storwize V7000 Gen1 compatibility mode cannot be reenabled.

Explanation: An attempt was made to enable Storwize V7000 Gen1 compatibility mode on a system where it was disabled. This mode cannot be reenabled.

User response: If you specified the wrong system, reenter the command. Otherwise, when a Storwize V7000 Gen1 enclosure is removed correctly, the data is migrated and the enclosure cannot be added back.

CMMVC8950E Storwize V7000 Gen1 compatibility mode cannot be disabled because it is not in use on this system.

Explanation: An attempt was made to disable Storwize V7000 Gen1 compatibility mode on a cluster for which it is not in use. Only Storwize V7000 systems (which can be hybrids) use Storwize V7000 Gen1 compatibility mode.

User response: If you specified the wrong system, retry the command with the correct system.

CMMVC8951E The Storwize V7000 Gen1 control enclosure cannot be added because Storwize V7000 Gen1 compatibility mode is not enabled.

Explanation: An attempt was made to add a Storwize V7000 Gen1 control enclosure to a system where Storwize V7000 Gen1 compatibility mode is disabled.

User response: Retry the command, and this time specify a Storwize V7000 Gen2 enclosure. Alternatively, add the Storwize V7000 Gen1 enclosure to a system where it is accepted.

CMMVC8952E The action failed because there was a mismatch with authentication credentials previously specified for same target.

Explanation: The **addiscsistorageport** command configures a session between an initiator port and target IP addresses for a target iSCSI Qualified Name (IQN). If you run the command again for the same target parameters, it returns successfully without initiating any action as the sessions was established after the first invocation. However, if you retry the command with a different **username** argument, **chapsecret** argument, or both, then the command fails because the initiator cannot verify which credentials are correct.

User response: To change the target credentials, complete the following steps by using the GUI or the specified command:

1. Use the **lsiscsistorageport** command to list the previous set of sessions that were established.
 2. Use the **rmiscsistorageport** command to remove the sessions that match the source port, IP address, and IQN for the target.
 3. Initiate a new discovery and establish sessions with new credentials.
-

CMMVC8953E There are no nodes in the site and/or I/O group specified.

Explanation: Nodes in the cluster were not configured to be part of the site before discovery of the iSCSI backend controller was initiated.

User response: If the system topology is a stretched cluster or HyperSwap configuration, use the management GUI or the **addnode** or **chnode** command to add the nodes to the specified site. Restart discovery.

CMMVC8954E The system is not in the replication layer.

Explanation: You can initiate discovery from a system in the replication layer only.

User response: Consider whether the source system is configured correctly. If appropriate, use the management GUI or the **chsystem -layer replication** command on the source system to configure the system in the replication layer. Rerun discovery by using the **detectiscsistorageportcandidate** command.

CMMVC8956E The command failed because the volume is the auxiliary volume in an active-active relationship.

Explanation: The specified volume is an auxiliary volume in an active-active Remote Copy relationship. You cannot throttle an auxiliary volume.

User response: Specify a different volume or break the Remote Copy relationship before you create or modify the throttle on this volume.

CMMVC8957E The command failed because the volume specified does not exist.

Explanation: A volume ID or name was specified that does not exist.

User response: View the list of volumes by means of the GUI or the **lsdisk** command to ensure that you specified the correct volume ID or name. For a HyperSwap volume involved in an active-active relationship, the volume ID and name are the same as those of the master volume.

CMMVC8958E The command failed because the volume has vvol ownership and may not have copies in two sites using HyperSwap.

Explanation: A volume created as a VVol may not be created as, or converted to, a HyperSwap volume that is involved in an active-active relationship.

User response: Use a different volume as a HyperSwap volume.

CMMVC8959E The command failed because the master volume has vvol ownership and may not participate in a Remote Copy relationship.

Explanation: A volume created as a VVol cannot

participate in a Remote Copy relationship as the master or auxiliary volume.

User response: Use VMware vCenter to manage system objects like volumes and pools.

CMMVC8960E The command failed because the specified node is the last node of the cluster.

Explanation: Only one node is left in the cluster. If you put this node into service state, the cluster is removed and cannot process the command. This error never occurs with **-action replace**.

User response: To put the last node of the cluster into service state, run **satask startservice**. This command does not run at the cluster level.

CMMVC8961E The action failed because the node hardware is incompatible with the previous node.

Explanation: A candidate node has the same World Wide Node Name as the node that you are trying to swap, but is a different node type, or has a different hardware configuration.

User response: Change your Fibre Channel configuration, and then retry the command.

CMMVC8962E The action failed because the specified node is not online.

Explanation: An attempt was made to run the **swapnode -action service** command on a node that is not online.

User response: First, make sure that you specified the correct node. If not, retry the action and specify an online node.

If you specified the correct node, then no further action is required because the node is already offline.

CMMVC8965E The command failed because volumes are dependent on the specified node and **-permitofflinevolumes** was not specified.

Explanation: Performing this operation may cause a host I/O outage. This error never occurs with **-action replace**.

User response: Correct the redundancy issue, or, if an I/O outage is acceptable, retry the command with the **-permitofflinevolumes** parameter.

CMMVC8966E This node cannot be added to the IO Group due to having lower configured RAM than required.

Explanation: An attempt was made to add a node to an I/O group where the other nodes in the group have more memory than the added node. It is also possible that faulty RAM is preventing the amount of memory in the node from being correctly detected. All nodes in an I/O group must have the same amount of configured RAM.

User response: Add the node to a different I/O group or increase its memory to match that of the other nodes in the group and retry the command.

CMMVC8967E The command failed because the auxiliary volume has vvol ownership and may not participate in a Remote Copy relationship.

Explanation: A volume created as a VVol cannot participate in a Remote Copy relationship as the master or auxiliary volume.

User response: Use VMware vCenter to manage system objects like volumes and pools.

CMMVC8968E This operation cannot be completed because iSCSI Initiator sessions are present.

Explanation: The operation is not permitted when iSCSI initiator sessions are present and you try to make one of the following changes:

- Changing from the replication layer to the storage layer on a system that is acting as an iSCSI initiator
- Changing from the storage layer to the replication layer on a system that is acting as an iSCSI target

User response: Follow these steps to remove the initiator sessions:

1. Use the management GUI or the **lsiscsistorageport** command to identify the iSCSI initiator sessions on the system that is acting as an iSCSI initiator.
2. Use the management GUI or the **rmiscsistorageport** command to remove all iSCSI initiator sessions.
3. After you remove all initiator sessions, retry the command.

CMMVC8969E This operation cannot be completed because migration to different parent storage pool is not supported.

Explanation: You cannot migrate volumes from a child storage pool to a different parent storage pool or to a child storage pool that has a different parent than the source.

User response: First migrate the volume from the child pool to its own parent, then migrate it from that parent pool to that of a different parent.

CMMVC8970E The storage pool was not deleted because there are volumes associated with the storage pool. A forced deletion is required.

Explanation: You cannot delete a storage pool that has any volumes still associated with it.

User response: Remove the associated volumes and retry the command, or use the **-force** flag:

```
rmmdiskgrp -force storage_pool
```

where *storage_pool* is either the name or the ID of a storage pool.

CAUTION:

Use of the -force flag can result in data loss. Contact IBM Support before using this flag.

CMMVC8971E 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: Check your command to ensure you have specified the correct volume and target. Make the correction and resubmit the command. Otherwise, investigate why the data has not been committed and how the data must be committed.

CMMVC8972E The command failed because the MDisk is an array.

Explanation: An attempt was made to change the encryption settings of an MDisk that is an array. The **chmdisk** command applies only to external MDisk.

User response: You cannot directly convert the array from unencrypted to encrypted or vice versa. Instead, you must complete one of the following actions:

- Specify a different MDisk and retry the command.
- Delete and re-create the array with the new encryption setting.

CMMVC8973W The IO group *io_group_name* has been restored with ID *new_id* instead of *old_id*.

Explanation: This situation can occur when the configuration node is different from the node that was used to create the original cluster. This change affects the SCSI Inquiry value for the I/O group.

User response: This message is a warning only. No user response is required.

CMMVC8974E The action failed because of incompatible code.

Explanation: The code version on one or more nodes is incompatible with the new version.

User response: Refer to the compatibility requirements for the code version you are adding. Update the cluster to meet the compatibility requirements, and then perform the upgrade.

CMMVC8975E The node could not be added because of incompatible code. The status code is *status_code*.

Explanation: An attempt was made to add a node with code that has a different revision level than the code in the cluster.

User response: Update the code on the rejected node to the same level as that on the cluster to which it will be added, and then resubmit the command.

CMMVC8976E 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.

CMMVC8977E 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/elog, /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 and try again.

CMMVC8978E 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: Check your command to ensure that you have the correct disk size. You can also check the **shrinkdisksize** command documentation for additional information. Make the correction and resubmit the command.

CMMVC8979E The action failed as the resulting disk size would be less than, or equal to, zero.

Explanation: An attempt was made to shrink a volume to a size that is below the permitted minimum.

User response: Check your command to ensure that you have the correct disk size. Make the correction and resubmit the command. For more information, see the documentation for the **shrinkvdisksize** command.

CMMVC8980E 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 refers to a disk that was destroyed.

User response: Correct the specified MDisk and resubmit the command.

CMMVC8981E The update failed as a file containing the code for the specified MCP version was not found.

Explanation: Two files are required to successfully complete a code update. One file contains the files that make up the base operating system, and the other file contains the code. This message appears if the OS version is incompatible with the code.

User response: Upload two compatible files and resubmit the command.

CMMVC8982E The action failed because the volume is part of a Remote Copy relationship.

Explanation: An action was performed against a volume that is part of a Remote Copy relationship.

User response: Remove the volume from the Remote Copy relationship and then resubmit the command.

CMMVC8983E The action failed because the volume 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.

CMMVC8984E 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.

CMMVC8985E The command has failed because a virtual medium error exists on the image mode volume 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.

CMMVC8986E 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.

CMMVC8987E 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: Inform your administrator of this error. You might need to update the code level on your server. Wait until the server is updated before you resubmit the command.

CMMVC8988E Failed to recreate the cluster you are trying to rebuild.

Explanation: An attempt was made to rebuild a cluster, but the attempt failed.

User response: Check your command. The source and target names might not match. Make the correction and resubmit the command.

CMMVC8989E 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 to another consistency group or remove an existing FlashCopy mapping from the desired group and then reissue the command.

CMMVC8990E The Remote Copy relationship was not created because the master or auxiliary volume is a member of a Remote Copy relationship.

Explanation: An attempt was made to create an active-active relationship with a volume copy in another site. This relationship is not supported if the volume is already in a remote copy relationship.

User response: Make sure that you specified the correct volume. If so, delete the existing relationship and try again.

CMMVC8991E 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.

CMMVC8992E 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.

CMMVC8993E The maximum number of WWPNS plus iSCSI qualified names (IQNs) for the cluster is already configured.

Explanation: The command cannot be initiated because the maximum number of WWPNS plus iSCSI qualified names for the cluster has been reached.

User response: Determine whether the action is required.

If the action is required, review the current configuration to determine whether any current WWPNS or iSCSI qualified name definitions are not required. Remove at least one WWPNS definition or iSCSI qualified name that is not required, and resubmit the command.

CMMVC8994E 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.

CMMVC8995E 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 WWPNS 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.

CMMVC8996E 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.

CMMVC8997E 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 and I/O group combination that is currently defined, and resubmit the command.

CMMVC8998E 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: Check your command. Change the host name if applicable and resubmit the command.

CMMVC8999E An IO group cannot be removed from a host because of one or more associated volumes.

Explanation: An I/O group cannot be removed from a host because of one or more associated VDisks.

User response: Check your command and ensure that you have specified the correct I/O group. Make the correction if one is needed and resubmit.

CMMVC9000E The action was not completed because the cluster has reached the maximum number of extents in storage pools.

Explanation: The cluster reached the maximum number of extents in the storage pool; therefore, the action did not complete. An attempt was made 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.

CMMVC9001I The package installed successfully.

Explanation: The package installed successfully.

User response: This message is for your information only. No response is required.

CMMVC9002E The cluster was recovered and the CLI functionality is limited until the cause of the failure is determined and any corrective action taken. Contact technical support for assistance.

Explanation: The cluster was recovered and the CLI functionality is limited.

User response: Contact IBM Support.

CMMVC9003E The action failed as the SSH key has been revoked.

Explanation: The action failed because the SSH key was revoked.

User response: Check your command and correct the number that you specified for the SSH key. Resubmit after you make the correction.

CMMVC9004E The action failed as the SSH key index (SSH_LABEL_ID) is invalid.

Explanation: The action failed because the SSH key index (SSH_LABEL_ID) is not valid.

User response: Correct the SSH key index and retry the command.

CMMVC6231E The action failed as the audit table is full.

Explanation: The action failed because the audit table is full.

User response: Save the audit log to disk and resubmit the command.

CMMVC9006E This operation cannot be performed because the cluster is currently cancelling the previous update command.

Explanation: This operation cannot be performed at the same time that the previous update command is being canceled.

User response: Wait until the previous update command stops running and resubmit the command.

CMMVC9007E This operation cannot be performed because, either an update has not been started, or an update is in progress but is not in a state where it can be aborted.

Explanation: This operation cannot be performed because the update is progressing.

User response: Wait until the update finishes and resubmit your command.

CMMVC9008E The update cannot be cancelled because at least one node has already committed to a new code level.

Explanation: You cannot cancel the update because one or more nodes already committed to a new code level.

User response: Contact IBM Support.

CMMVC9009E 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 was entered. The command was not run.

User response: Enter either yes or no.

CMMVC9010E The command has not completed. A limited availability parameter has been entered without the required environment setting being set.

Explanation: The command did not finish. A limited availability parameter was entered without the required environment setting.

User response: Check your command. If the command is correct, contact your administrator to determine whether you need an environment setting so that the command can run. Resubmit the command when the problem is resolved.

CMMVC9011E The command failed as the remote cluster does not support global mirror.

Explanation: The command failed because the remote cluster does not support Global Mirror.

User response: Check your command and ensure that you specified the correct cluster. Make the correction and resubmit the command. If you specified the correct cluster, research why it does not support Global Mirror.

CMMVC9012E 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.

CMMVC9013E 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.

CMMVC9014E 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.

CMMVC9015E 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.

CMMVC9016E 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 you start it.

CMMVC9017E 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 was already submitted and is still in progress. When the task completes successfully, the FlashCopy mapping or consistency group state changes to Stopped.

User response: Wait for the existing task to complete. The group state changes to Stopped automatically.

CMMVC9018E The FlashCopy mapping was not created because the source volume 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. A source volume was specified that is currently defined as the target of a FlashCopy mapping.

User response: You have two options. One option is to specify a different source volume and resubmit the command. The other option is to delete the existing FlashCopy mapping that defines the source volume that was specified as the target volume, and resubmit the command.

CMMVC9019E 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 was 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 to specify a different source volume and resubmit the command. The other option is to delete one of the existing FlashCopy mappings that contains the source volume and resubmit the command.

CMMVC9020E The FlashCopy mapping was not created because the target volume is already a source volume 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 was specified is currently defined as the source of a FlashCopy mapping.

User response: You have two options. One option is to specify a different target volume and resubmit the

command. The other option is to delete all of the existing FlashCopy mappings that contain the target volume that was specified and resubmit the command.

CMMVC9021E 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 was specified is currently defined as the target of another FlashCopy mapping.

User response: You have two options. One option is to specify a different target volume and resubmit the command. The other option is to delete the existing FlashCopy mapping that contains the target volume that was specified and resubmit the command.

CMMVC9022E The command failed because the authorization table is full.

Explanation: The command failed because the authorization table is full.

User response: Check with your administrator to verify the status of the authorization table. You must wait until the table is adjusted to resubmit the command.

CMMVC9023E 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: Check your command for the role you specified. If the role is not set up, check with your administrator to verify the status. Resubmit the command after the role is set up in the system.

If you want to set the role as the default, no further action is required.

CMMVC9024E 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 the **rmauth** command to set the default role.

CMMVC9025E The command failed because the specified role was not found.

Explanation: The command failed because the specified role was not found.

User response: Check the command and correct the specified role. If you believe the specified role exists, check with your administrator for clarification. Make the correction and resubmit the command.

CMMVC9026E 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: Check the command and ensure that you specified a valid SSH key. Make the correction and resubmit the command.

CMMVC9027E 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.

CMMVC9028E 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: Recheck your command and ensure that you specified admin as the SSH user name. Make the correction and resubmit.

CMMVC9029E The command cannot set the authorization record to the default role.
Use **rmauth** to set the default role.

Explanation: The command cannot set the authorization record to the default role.

User response: Use the **rmauth** command to set the default role.

CMMVC9030E The command failed because the SSH key already exists or there is a duplicate SSH key.

Explanation: An attempt was made to add an SSH key that already exists. A different authorization level might be associated with the key.

User response: If the existing SSH key of the same type does not have the authority level that you require, add a different SSH key.

CMMVC9031E The command failed because one of the nodes in the specified I/O group was offline when a memory change was attempted.

Explanation: All nodes in the I/O group must be online when you enter the **chiogrp** command. This error indicates that at least one node is not currently online.

User response: Check the status of the nodes by using the **lsnodecanister** command for enclosure-based systems or the **lsnode** command for appliance-based systems. If any node in the specified I/O group is not currently online, follow maintenance procedures for that node. When all nodes are online, repeat the command.

CMMVC9032E The Add E-mail User operation failed because there is no space left in the user list.

Explanation: The maximum number of email recipients is already configured.

User response: Remove an existing email recipient and retry.

CMMVC9033E The operation failed because there is already a user with that name.

Explanation: An email recipient with that address already exists.

User response: Make sure that you specified the correct user name. If not, make the correction and try again.

CMMVC9034E The operation failed because the specified user does not exist.

Explanation: An attempt was made to perform an operation on a user that does not exist.

User response: Retry the command with an existing user.

CMMVC9035E The Remove E-mail User operation failed because this is the last entry in the user list.

Explanation: Email services require that at least one participant is configured.

User response: Run the **stopemail** command to stop email services and then remove the email user.

CMMVC9036E Sendmail error EX_USAGE. A command or configuration line has been used incorrectly.

Explanation: The send email task failed because a command or a configuration line was used incorrectly.

User response: Ensure that the email settings are correct and resubmit the task.

CMMVC9037E Sendmail error EX_DATAERR. Address is wrong, or the message is too large for the mailbox.

Explanation: The send email task 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.

CMMVC9038E Sendmail error EX_NOINPUT. An input file (not a system file) did not exist or was not readable.

Explanation: The send email task failed because a file is missing or cannot be read.

User response: Ensure that the email system is configured correctly. Ensure that access permissions are specified correctly for all email configuration files. Resubmit the task.

CMMVC9039E Sendmail error EX_NOUSER. The sendmail command could not recognize a specified user.

Explanation: The send email task 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.

CMMVC9040E Sendmail error EX_NOHOST. The sendmail command could not recognize the specified host name.

Explanation: The send email task failed because the host is not known to the email system.

User response: Ensure that you configured the SMTP

environment correctly and that you specify a defined host. Resubmit the task.

CMMVC9041E Sendmail error EX_UNAVAILABLE. A required system resource is not available.

Explanation: The send email task failed because a required system resource is not available.

User response: Ensure that you configured the SMTP environment correctly, and then resubmit the task.

CMMVC9042E Sendmail error EX_SOFTWARE. An internal error occurred (including bad arguments).

Explanation: The send email task failed because an incorrect parameter or parameter value was detected.

User response: Ensure that you configured the SMTP environment correctly. Specify only supported parameters and parameter values, and resubmit the task.

CMMVC9043E Sendmail error EX_OSERR. A system resource error prevented the sending of an email.

Explanation: The send email task failed because a system resource error occurred.

User response: Ensure that you configured the SMTP environment correctly and resubmit the task.

CMMVC9044E Sendmail error EX_OSFILE. Failed to open a critical system file.

Explanation: The send email task failed because a required system file cannot be opened.

User response: Ensure that the email system is configured correctly, and that access permissions are specified correctly for all email configuration files. Resubmit the task.

CMMVC9045E Sendmail error EX_CANTCREAT. An output file could not be written to by sendmail.

Explanation: The send email task failed because the system cannot write to a required output file.

User response: Ensure that the email system is configured correctly, and that access permissions are specified correctly for all email configuration files. Resubmit the task.

CMMVC9046E 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 failed because a write or read I/O operation 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.

CMMVC9047E Sendmail error EX_TEMPFAIL. The sendmail command could not create a connection to a remote system.

Explanation: The send email task 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.

CMMVC9048E Sendmail error EX_PROTOCOL. The remote system returned something that was incorrect during a protocol exchange.

Explanation: The send email task failed because an error occurred in the protocol exchange.

User response: Ensure that the email system is configured correctly and that you configured the SMTP environment correctly. Resubmit the task.

CMMVC9049E Sendmail error EX_NOPERM. The user does not have permission to perform the requested operation.

Explanation: The send email task 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.

CMMVC9050E Sendmail error EX_CONFIG. There is a fatal problem with the sendmail configuration.

Explanation: The send email task failed because the sendmail configuration is not correct.

User response: Ensure that the email system is configured correctly and that you configured the SMTP environment correctly. Resubmit the task.

CMMVC9051E An unknown error occurred. Ensure your SMTP server is running.

Explanation: The send email task failed because an unexpected error occurred.

User response: Ensure that the SMTP server is running, and resubmit the task.

CMMVC9052E The email command timed out. Check your email server settings as listed on the cluster.

Explanation: The send email task failed because a command timeout occurred.

User response: Ensure that your system settings match those that are recommended in the sendmail application documentation, and resubmit the task.

CMMVC9053E The email service has not been enabled.

Explanation: The send email task failed because the email application is not enabled.

User response: Enable the email application by using the **startemail** command and resubmit the task.

CMMVC9054E 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.

CMMVC9055E The command failed because a target volume 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 depend 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.

CMMVC9056E The create failed because the source and target volumes 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. An attempt was made to create a new FlashCopy mapping that links two existing trees that have different grain sizes.

User response: You have three options.

- Resubmit the command and specify a different source or target volume.
- Delete all of the existing mappings that contain the source volume and resubmit the command.
- Delete all of the existing mappings that contain the target volume and resubmit the command.

CMMVC9061E Cannot disable key server type because key server objects of this type exist.

Explanation: An attempt was made to disable a key server type for which a key server end point exists. The end point indicates that the key server type is actively being used for encryption and therefore cannot be disabled.

User response: Verify that you specified the correct key server type. If so, delete all key server objects of this type by using the **rmkeyserver** command and try again.

CMMVC9062E Cannot disable key server type because it is not currently enabled.

Explanation: An attempt was made to disable a key server type that is not currently enabled. The **-disable** parameter can only be used to disable a key server type that is currently enabled and has no online key server objects.

User response: Retry the command and specify an appropriate key server type, that is, one that is currently enabled and has no online key server objects.

CMMVC9064E This host is the only host in the host cluster. The host cluster will lose all of its shared mappings after removing this host. Use -force flag to continue.

Explanation: An attempt was made to remove the last host in a host cluster that still includes shared mappings. You must use the **-force** flag in these circumstances, which removes all shared mappings for the host cluster.

User response: If you are sure that you want to delete the host, retry the command and use the **-force** flag.

Note: Use of the **-force** flag can result in unintended loss of data.

CMMVC9065E Host cluster does not contain any hosts.

Explanation: An attempt was made to carry out one of the following actions on an empty host cluster:

- Modify the cluster
- Add or remove shared mappings

User response: Make sure that you specified the correct host cluster. If so, add hosts to the host cluster before you modify it, or before you add or remove shared mappings.

CMMVC9066E Volume already has a shared mapping to the host cluster.

Explanation: An attempt was made to map a volume to the same host cluster twice.

User response: Correct the volume name and retry the command.

CMMVC9067E Volume is mapped as a subsidiary LUN (Virtual Volume).

Explanation: An attempt was made to map a VMware vSphere Virtual Volume to a host cluster. This mapping is not permitted.

User response: Retry the command and use a volume that is not a Virtual Volume.

CMMVC9068E Hosts in the host cluster have conflicting SCSI ID's for their private mappings.

Explanation: An attempt was made to create a shared host cluster mapping to a volume, but a host within the host cluster is already privately mapped to that volume.

User response: Make sure that the host mappings are compatible with the mappings for the host cluster and the hosts it contains.

CMMVC9069E Volume does not have a shared mapping to this host cluster.

Explanation: An attempt was made to remove a mapping from a volume to a host cluster where the mapping did not exist.

User response: Make sure that you specified the correct volume and host cluster. If not, retry the command with the correct parameters. If so, no further action is necessary because the mapping does not exist.

CMMVC9070E Either -keepmappings or -removemappings flag must be specified.

Explanation: An attempt was made to remove a host from a host cluster that has one or more shared mappings, but no flag was set to determine the disposition of the mappings for that host.

User response: Retry the command and specify either the **-keepmappings** flag to retain the existing mappings or the **-removemappings** flag to delete them.

CMMVC9071E Cannot remove a mapping because it is a shared mapping.

Explanation: An attempt was made to remove a shared mapping from a host cluster with the **rmvdiskhostmap** command, which removes a single host-volume mapping.

User response: Use the **rmvolumehostclustermap** command to remove a shared mapping from a host cluster.

CMMVC9072E The host cluster mapping was not created because a volume is already mapped to this host cluster with this SCSI LUN.

Explanation: A SCSI LUN was specified that was already in use for a volume that is mapped to the current host cluster.

User response: Retry the command, specifying a different SCSI LUN.

CMMVC9073E The host cluster mapping was not created because this volume is mapped to one of the hosts in the host cluster with a different SCSI LUN.

Explanation: An incompatible SCSI LUN ID was specified for the mapping. Hosts within the host cluster have private mappings to volumes with this SCSI LUN.

User response: Retry the command and specify a compatible SCSI LUN ID, that is, an ID that is not in use for private mappings.

CMMVC9074E The host cluster mapping was not created because this volume is mapped to a host not in this host cluster. Use -force to create the mapping.

Explanation: An attempt was made to map a volume to a host cluster when that volume is already mapped to another host or host cluster.

User response: Make sure that you specified the correct volume and host cluster. If so, you can retry the command and use the **-force** flag to create the mapping. If you do so, the previous mapping will also be preserved. Always be cautious when you use the **-force** flag, which can have unforeseen consequences.

CMMVC9075E The host cluster mapping was not created because the limit of SCSI LUNs supported for this host type has been reached.

Explanation: An attempt was made to change the host cluster type to **adminlun**, but one of the mappings for one of the hosts in the host cluster has a SCSI LUN that exceeds the maximum for the **adminlun** type. The maximum number of SCSI LUNs for the default host

type is 4096 (2048 for older hardware), and the maximum for the **adminlun** host type is 512.

User response: Reduce the number of SCSI LUNs by using the **rmvdiskhostmap** command to remove at least one host mapping, then retry the command.

CMMVC9076E The host cluster mapping was not created because no free SCSI LUN could be found for this shared mapping.

Explanation: An attempt was made to create a shared mapping, but no free SCSI LUN slot was found that was available for all of the hosts in the host cluster.

User response: Unmap one or more volumes to free up at least one SCSI LUN slot and resubmit the command.

CMMVC9077E The host cluster mapping was not created because of conflicting SCSI LUNs. One of the hosts in the host cluster already has a mapping to this volume, however, another host in the host cluster has a mapping to a different volume but with the same SCSI LUN.

Explanation: An attempt was made to create a shared mapping, but hosts in the host cluster have mappings with the same SCSI LUN but different volumes.

User response: Remap the volume to a compatible SCSI LUN ID before you change to a shared mapping, or remove the private mappings for other hosts to eliminate the conflicts.

CMMVC9078E A volume listed in the -ignoreseedvolume list does not have a mapping to any of the hosts specified in the -seedfromhost list.

Explanation: A volume was specified as "non-seeding," but that volume is not mapped to the seeding host.

User response: Retry the command and specify a volume that is mapped to the seeding host.

CMMVC9079E Cannot delete all the hosts, because at least one of the hosts has mappings. After deleting all the hosts, all of their mappings will be deleted as well. Use -force flag to continue.

Explanation: An attempt was made to remove multiple hosts from a host cluster while at least one of the hosts was still mapped to a volume.

User response:

- Delete the mappings for any hosts that you want to remove from a host cluster.

- Alternatively, use the **-force** flag to automatically delete all the mappings for a host when that host is removed.

Note: Use of the **-force** flag can have unintended consequences.

CMMVC9080E The host cluster I/O group(s) cannot be removed because at least one volume in the I/O group(s) has received I/O within the defined volume protection period.

Explanation: An attempt was made to remove an I/O group where global volume protection is enabled and I/O occurred within the specified quiet period.

User response: Make sure that you specified the correct I/O group. If so, make sure that no activity that involves the I/O group takes place during the required quiet period before you retry the command.

CMMVC9081E Duplicate host IDs were entered.

Explanation: An attempt was made to add or remove multiple hosts from a host cluster where the same host ID was specified twice.

User response: Correct the host IDs, specifying each ID only once, and retry the command.

CMMVC9082E Too many hosts supplied in the list.

Explanation: A list of hosts that are to be added to a host cluster was provided, but the number of hosts in the list is too large. A host cluster can include a maximum of 128 hosts.

User response: Retry the command with a shorter list of hosts.

CMMVC9083E A new cloud snapshot could not be created because the volume is not ready.

Explanation: A cloud snapshot cannot be created if any of the following conditions apply:

- A cloud snapshot, restore, or delete operation is already in progress on the volume
- An unfixed cloud snapshot error is logged against the volume

A new cloud snapshot can be started only when the volume **backup_status** is **ready**.

User response: Complete the following actions:

- Fix any outstanding snapshot errors and retry the command.
- If the error persists, use the management GUI or the **lsvolumebackupprogress** and **lsvolumerestoreprogress** commands to monitor the progress of existing cloud snapshot, delete, and

restore operations. Wait for the existing operation to complete, or cancel the operation, and then retry the original command.

CMMVC9084W Volumes that are enabled for cloud snapshots exist without a license for each enclosure.

Explanation: Each enclosure that contains a volume that is enabled for cloud snapshots must have a valid transparent cloud tiering license. At least one enclosure is missing a license.

User response: Obtain a valid transparent cloud tiering license for any enclosures that require one.

CMMVC9085E The command failed because the cloud account is still being initialized.

Explanation: Initialization of the cloud account must complete before the cloud account can be used.

User response: Wait for initialization of the cloud account to complete, then retry the command.

CMMVC9086E A new cloud snapshot could not be created because the cloud account configured for the volume is in import mode.

Explanation: A new cloud snapshot can be created only when the cloud account is in normal mode.

User response: Change the mode of the cloud account or configure the volume to use a different cloud account.

CMMVC9087E A new cloud snapshot cannot be created because the volume is not enabled for cloud snapshots.

Explanation: A new cloud snapshot can be created only when the cloud snapshot function is enabled.

User response: Enable the cloud snapshot function for the volume by using the management GUI or the **chvdisk** command.

CMMVC9088E A new cloud snapshot could not be created because the maximum number of cloud snapshots already exists for the volume.

Explanation: The volume already has the maximum number of cloud snapshots.

User response: Remove any unwanted cloud snapshots for the volume and retry the command. Use the management GUI or the **lsvolumebackupgeneration** and **rmvolumebackupgeneration** commands to list and remove cloud snapshots.

CMMVC9089E The command failed because no cloud snapshots for the specified volume exist.

Explanation: An attempt was made to remove cloud snapshots where a valid volume ID was provided, but no cloud snapshots for the specified volume were found.

User response: Use the management GUI or the **lsvolumebackup** command to list the volumes for which cloud snapshots exist in the cloud. Specify one of these volumes when you retry the command.

CMMVC9090E The command failed because the specified cloud snapshot does not exist.

Explanation: An attempt was made to remove or restore a cloud snapshot, but the specified snapshot was not found in the cloud.

User response: Use the management GUI or the **lsvolumebackupgeneration** or **lsvolumebackup** command to list the cloud snapshots that exist in the cloud for a particular volume. Specify one of these snapshots when you retry the command.

CMMVC9091E The command failed because an existing delete operation is in progress for this volume.

Explanation: Only one delete operation is allowed at a time for a volume that has cloud snapshots. The delete process is asynchronous and runs in the background. While that process is running for one snapshot, you cannot delete another snapshot of the same volume.

User response: Use the management GUI or the **lsvolumebackupprogress** command to list the progress of current delete operations. When no other delete operations are in progress for the volume, retry the command.

CMMVC9092E The command failed because the specified cloud snapshot is the most recent complete cloud snapshot for the volume.

Explanation: An attempt was made to delete the most recent cloud snapshot for a volume. The most recent snapshot must remain available for use in restore operations.

User response: Use the management GUI or the **lsvolumebackup** to choose a different cloud snapshot to delete.

CMMVC9093E The command failed because a restore operation is in progress from this cloud snapshot.

Explanation: An attempt was made to delete a cloud snapshot while it was being used to restore a volume.

User response: Wait for the restore operation to complete, or cancel the restore operation. Use the management GUI or the **lsvolumerestoreprogress** command to display the progress of current restore operations.

CMMVC9094E The command failed because a restore operation is in progress from one of the cloud snapshots for this volume.

Explanation: An attempt was made to delete all cloud snapshots for a volume while one of the cloud snapshots is being used to restore a volume.

User response: Wait for the restore operation to complete, or cancel the restore operation. Use the management GUI or the **lsvolumerestoreprogress** command to list the progress of current restore operations.

CMMVC9095E The command failed because a volume with the UID specified with the **-fromuid** parameter already exists on the local system.

Explanation: The **-fromuid** parameter is used to restore a cloud snapshot that was taken from a different volume.

User response: Do not use the **--fromuid** parameter when the specified volume already exists on the local system.

CMMVC9096E The command failed because an existing restore operation is in progress on this volume.

Explanation: An attempt was made to restore a volume that was already being restored.

User response: Wait for the restore operation to complete. If you specified the wrong volume, you can display a list of all volumes for which a restore operation is in progress by using the management GUI or the **lsvolumerestoreprogress** command. Specify a volume for which no restore operations are in progress when you retry the original command.

CMMVC9097E The command failed because the specified version of the cloud snapshot is not the latest cloud snapshot for the volume. The **-deletelatergenerations** parameter must be specified if the latest version of the cloud snapshot is not specified.

Explanation: If cloud snapshots are enabled on the volume, and the cloud snapshot that is being restored is not the most recent cloud snapshot of the volume, then all later cloud snapshots for the volume must be deleted by specifying the **-deletelatergenerations** parameter.

User response: Complete one of the following actions:

- Retry the command with the **-deletelatergenerations** parameter.
- If you are restoring directly to the production volume, disable cloud snapshots on the volume and retry the command. If cloud snapshots are enabled later, the next cloud snapshot will be a full cloud snapshot.
- If you are committing a restore from a temporary volume, consider converting the temporary volume into an independent volume by using the **-detach** parameter. You can then keep the original volume and volume cloud snapshots unchanged.

CMMVC9098E The command failed because an existing restore operation is in progress from this volume cloud snapshot.

Explanation: An attempt was made to restore a volume where the specified cloud snapshot was already in use for a restore operation.

User response: Use the management GUI or the **lsvolumerestoreprogress** command to list the progress of current restore operations.

CMMVC9099E The command failed because no restore operation is in progress for this volume.

Explanation: An attempt was made to cancel a restore for a volume for which no restore operation is in progress.

User response: To display a list of all volumes for which a restore operation is in progress, use the management GUI or the **lsvolumerestoreprogress** command. Specify one of these volumes when you retry the original command.

CMMVC9100E The command failed because there is no uncommitted restore operation for this volume.

Explanation: An attempt was made to commit or detach a restore for a volume for which no uncommitted restore operation was found.

User response: To display a list of all volumes for which a restore operation is in progress, use the management GUI or the **lsvolumerestoreprogress** command. Specify one of these volumes when you retry the original command.

CMMVC9101E The command failed because the -fromuid parameter was specified and the specified local volume has cloud snapshot enabled.

Explanation: An attempt was made to restore a cloud snapshot from a different volume. However, cloud

snapshots are enabled on the specified local volume.

When you use the **-fromuid** parameter, cloud snapshots cannot be enabled on the local volume.

User response: Complete one of the following actions:

- Disable cloud snapshots on the local volume and retry the command.
- Restore to a different local volume.

CMMVC9102E The command failed because the -restoreuid parameter was specified when the existing volume has mappings to hosts.

Explanation: An attempt was made to restore a cloud snapshot from a different volume, where the UID of the local volume was requested to be set to the UID of the volume cloud snapshot. However, the local volume has mappings to a host object.

When you use the **-restoreuid** parameter, the local volume cannot have any volume to host mappings.

User response: Complete one of the following actions:

- Remove the volume to host mappings for the local volume and retry the command.
- Restore to a different local volume.

CMMVC9103E A restore operation could not be started because the volume is not ready.

Explanation: A restore operation cannot be started if any of the following conditions apply:

- A cloud snapshot, restore, or delete operation is already in progress on the volume
- An unfixed cloud snapshot error is logged against the volume

A restore operation can be started only when the volume **restore_status** is **available**.

User response: Complete the following actions:

1. Fix any outstanding snapshot errors and retry the command.
2. If the error persists, use the management GUI or the **lsvolumebackupprogress** and **lsvolumerestoreprogress** commands to monitor the progress of existing cloud snapshot, delete, and restore operations. Wait for the existing operation to complete, or cancel the operation, and then retry the original command.

CMMVC9104E A cloud snapshot could not be deleted because the volume is not ready.

Explanation: A cloud snapshot cannot be deleted if any of the following conditions apply:

- A cloud snapshot, restore, or delete operation is already in progress on the volume

- An unfixed cloud snapshot error is logged against the volume

User response: Complete the following actions:

1. Fix any outstanding snapshot errors and retry the command.
2. If the error persists, use the management GUI or the **lsvolumebackupprogress** and **lsvolumerestoreprogress** commands to monitor the progress of existing cloud snapshot, delete, and restore operations. Wait for the existing operation to complete, or cancel the operation, and then retry the original command.

CMMVC9105E The command failed because the local volume is a different size than the specified cloud snapshot.

Explanation: An attempt was made to restore a volume that does not have the same virtual capacity as the volume cloud snapshot.

User response: Resize the local volume and retry the command, or restore to an alternative local volume.

CMMVC9106E The command failed because the **-createtemporaryvolume** option can be used only if the volume is enabled for cloud snapshots.

Explanation: An attempt was made to create a temporary volume to hold the restored data, but cloud snapshots are not enabled on the target volume for the restore operation. This usage is not supported.

User response: Complete one of the following actions:

- Run the restore operation directly to the specified volume by omitting the **-createtemporaryvolume** parameter.
- Enable cloud snapshots on the target volume for the restore operation and retry the command.

CMMVC9107E A temporary volume could not be created for the restore process because there are insufficient free volume IDs or volume copy IDs in the system.

Explanation: An attempt was made to create a temporary volume to hold the restored data, which required that a new volume be created in the same caching I/O group as the target volume for the restore operation. The new volume was not created because no free volume IDs or volume copy IDs were available in the system.

User response: Determine whether the system includes unwanted volumes or volume copies. If so, remove unwanted volumes or volume copies by using the **rmvolume** or **rmvolume copy** command, and then retry the command.

CMMVC9108E The volume cannot be resized because the volume is enabled for cloud snapshots.

Explanation: An attempt was made to change the size of a volume when the cloud snapshot feature was enabled on the volume.

User response: Make sure that you specified the correct volume. If so, disable the cloud snapshot feature on the volume by using the **chvdisk** command and retry the original command.

CMMVC9109E A temporary volume could not be created for the restore process because there is insufficient free capacity in the storage pool.

Explanation: An attempt was made to create a temporary volume to hold the restored data, which required that a new volume be created in the same caching I/O group as the target volume for the restore operation. The new volume was not created because no free volume IDs were available in the storage pool.

User response: Determine whether the system includes unwanted local volumes. If so, remove unwanted local volumes by using the **rmvdisk** command, and then retry the command.

CMMVC9110E The restore operation could not be started because the maximum number of FlashCopy mappings in the system has been reached.

Explanation: The restore operation requires an internal FlashCopy mapping to be created and the system limit has been reached.

User response: Remove unwanted FlashCopy mappings from the system by using the **rmfcmap** command and retry the **restorevolume** command.

CMMVC9111E The command failed because the specified volume is part of a Metro Mirror or Global Mirror relationship.

Explanation: An attempt was made to restore to a volume that is part of a remote copy relationship.

User response: Make sure that you specified the correct volume. If so, remove the remote copy relationship that uses this volume by using the **rmrcrelationship** command and retry the **restorevolume** command.

CMMVC9112E The command failed because the specified volume is a HyperSwap volume.

Explanation: An attempt was made to restore to a HyperSwap volume. This operation is not allowed.

User response: Make sure that you specified the correct volume. If so, convert the volume to a basic volume by removing a copy at one site, then retry the command.

CMMVC9113E Cloud snapshots are already enabled for the specified volume.

Explanation: An attempt was made to enable cloud snapshots on a volume where cloud snapshots are already enabled.

User response: Make sure that you specified the correct volume. If you want to enable cloud snapshots that use a different cloud account, first disable cloud snapshots for the current account, then retry the command.

CMMVC9114E Cloud snapshots cannot be disabled because they are not enabled on this volume.

Explanation: An attempt was made to disable the cloud snapshot function when it was not enabled.

User response: Make sure that you specified the correct volume.

CMMVC9115E Cloud snapshots cannot be enabled because the maximum number of cloud snapshot-enabled volumes in the system has been reached.

Explanation: The system limit was reached for the number of volumes that can be enabled for cloud snapshots.

User response: You must disable cloud snapshots for another volume before you can enable cloud snapshots for the current volume.

CMMVC9116E Cloud snapshots cannot be enabled because the specified cloud account is in import mode.

Explanation: The cloud account must be in normal mode when you enable cloud snapshots.

User response: Change the account to normal mode and retry the command.

CMMVC9117E The volume was not deleted because a cloud snapshot operation is in progress.

Explanation: The volume was not deleted because it would prevent a cloud snapshot operation from completing.

User response: Make sure that you specified the correct volume. If so, wait for the cloud snapshot operation to complete, or cancel the cloud snapshot. Use the management GUI or the

lsvolumebackupprogress command to monitor the progress of the cloud snapshot, or the **rmvolumebackupgeneration** command to cancel the cloud snapshot. Alternatively, specify the **-cancelbackup** parameter to force-delete the volume, canceling any active cloud snapshot operations on the volume.

CMMVC9118E The volume was not deleted because a restore operation is in progress and it would allow access to an image mode copy that does not contain consistent data.

Explanation: The volume was not deleted because an image mode copy might contain inconsistent data.

User response: Make sure that you specified the correct volume. If so, wait for the restore operation to complete, or cancel the restore. Use the management GUI or the **lsvolumerestoreprogress** command to monitor the progress of the restore, or the **restorevolume** command to cancel the restore. Alternatively, specify the **-discardimage** parameter to force-delete the volume, canceling any active restore operation.

CMMVC9119E The volume copy was not deleted because a restore operation is in progress and it might allow access to an image mode copy that does not contain consistent data.

Explanation: The volume copy was not deleted because an image mode copy might contain inconsistent data.

User response: Make sure that you specified the correct volume. If so, wait for the restore operation to complete, or cancel the restore. Use the management GUI or the **lsvolumerestoreprogress** command to monitor the progress of the restore operation, or the **restorevolume** command to cancel the restore operation. Alternatively, specify the **-discardimage** parameter to force-delete the volume copy.

CMMVC9120E The remote copy relationship was not created because the master or auxiliary volume is enabled for cloud snapshots.

Explanation: A volume where cloud snapshots are enabled cannot be part of a remote copy relationship.

User response: Make sure that you specified the correct volume. If so, disable the cloud snapshot feature by using the **rmcloudaccount** command and retry the original command.

CMMVC9121E The change volume could not be associated because the volume is enabled for cloud snapshots.

Explanation: A volume cannot be configured as a change volume for a remote copy relationship if the cloud snapshot feature is enabled on the volume.

User response: Make sure that you specified the correct volume. If so, disable the cloud snapshot feature on the volume by using the **rmcloudaccount** command and retry the original command.

CMMVC9122E The FlashCopy mapping was not created because the source or target volume is enabled for cloud snapshots.

Explanation: A volume cannot be the source or target volume in a FlashCopy mapping if the cloud snapshot feature is enabled on the volume.

User response: Make sure that you specified the correct volume. If so, disable the cloud snapshot feature on the volume by using the **chvdisk** command and retry the original command.

CMMVC9123E The command failed because a key server type has not been enabled.

Explanation: An attempt was made to create a key server without enabling a key server type.

User response: Enable a key server type by using the **chkeyserverisklm** command. Then, retry the **mkkeyserver** command.

CMMVC9124E The command failed because a primary ISKLM key server already exists.

Explanation: An attempt was made to create a key server object by using the **mkkeyserver -primary** command where the primary ISKLM key server already exists. After the primary key server is created, subsequent key server objects cannot be created with the **-primary** flag. The primary ISKLM key server can be changed later by using the **chkeyserver** command.

User response: Retry the command without specifying the **-primary** flag.

CMMVC9125E The command failed because the **-primary** flag must be set when creating the first ISKLM key server.

Explanation: An attempt was made to create the primary ISKLM key server object without specifying the **-primary** flag. The first key server that you create of the ISKLM type must have the **-primary** flag set. In other words, you must create your primary key server first. Subsequent key server objects cannot be created with the **-primary** flag. The primary ISKLM key server

can be changed later by using the **chkeyserver** command.

User response: Specify the **-primary** flag when you retry the command.

CMMVC9126E The command failed because the **-primary** flag can only be set for ISKLM key servers.

Explanation: An attempt was made to create a key server object where the **-primary** flag was specified. The **-primary** flag is valid only for ISKLM key servers, but the ISKLM key server type is not currently enabled.

User response: Retry the command without specifying the **-primary** flag.

CMMVC9127E The command failed because an SSL certificate is required.

Explanation: An attempt was made to create a key server object without specifying a self-signed SSL certificate. The parent key server type object does not have a CA certificate configured either. As a result, no SSL certificate would be available on the system for communicating with the key server.

User response: Configure a CA certificate for this key server type or supply a self-signed certificate for the key server with the **-sslcert** parameter.

CMMVC9128E Cannot enable key server type because it would exceed the permitted number of enabled key server types.

Explanation: One key server type is typically enabled at a time. Migration between key server types is a special case where a second type can be enabled until the migration is complete. This error means that an attempt was made to enable a third key server type during a migration, which is not permitted.

User response: If you are currently migrating between two key server types, try the command again and this time include the **-disable** parameter to disable the migration target.

CMMVC9129E The command has failed because an IPv4 address was specified and each node does not have an IPv4 service IP address.

Explanation: An attempt was made to create the key server object with an IPv4 address, which requires each node in the system to have an IPv4 service IP address set.

User response: Use the **lsservicestatus** command to ensure that each node in the system has an IPv4 service IP address and then retry the command. Alternatively, if each node has an IPv6 service address, specify an IPv6 service IP address when you retry the command.

CMMVC9130E The command has failed because an IPv6 address was specified and each node does not have an IPv6 service IP address.

Explanation: An attempt was made to create the key server object with an IPv6 address, which requires each node in the system to have an IPv6 service IP address set.

User response: Use the **lsservicestatus** command to ensure that each node in the system has an IPv6 service IP address and then retry the command. Alternatively, if each node has an IPv4 service address, specify an IPv4 service IP address when you retry the command.

CMMVC9131E The command failed because the key server reported an error.

Explanation: An attempt was made to create or test a key server on the system. During both the **mkkeyserver** and **testkeyserver** tasks, the system attempts to validate the key server. During this validation, the key server reported an error. More detailed information about the error might be found in the additional sense data in the event log.

User response: Check the event log for key server errors. Fix any key server errors and run the task again.

CMMVC9132E This will change the site of at least one host in the host cluster that already has a site defined. Use **-force** flag to continue.

Explanation: An attempt was made to change the site of the host cluster, which would change the site of at least one host in the host cluster.

User response: Make sure you specified the correct information in the command. If so, you can use the **-force** flag to forcibly change the site of one or more hosts in the host cluster. Using the **-force** flag can have unexpected results.

CMMVC9134E The command failed because the key server is not supported.

Explanation: An attempt was made to create or test a key server on the system. During both the **mkkeyserver** and **testkeyserver** tasks, the system attempts to validate the key server. During this validation, the key server reported unsupported vendor information. More detailed information about the server might be found in the additional sense data in the event log.

User response: Check the event log for key server errors. Fix any key server errors and run the task again.

CMMVC9135E The command failed because there was a problem establishing a connection to the key server.

Explanation: An attempt was made to create or test a key server on the system. During both the **mkkeyserver** and **testkeyserver** tasks, the system attempts to validate the key server. An error occurred in establishing a connection to the key server by using the supplied IP address, IP port, and SSL certificate for the key server. This error might be caused by a network problem, incorrect IP address or port details, or a problem with the SSL certificate. More detailed information about the error might be found in the additional sense data in the event log.

User response: Confirm that the correct IP address, IP port, and SSL certificate were supplied for the key server. Confirm that each node in the system has access to the key server. Confirm that the key server is fully operational and run the task again.

CMMVC9136E The command failed because the key server's response was not understood.

Explanation: An attempt was made to create or test a key server on the system. During both the **mkkeyserver** and **testkeyserver** tasks, the system attempts to validate the key server. The system was unable to process the response from the key server.

User response: Fix any key server errors and run the task again.

CMMVC9137E The command failed because no primary key server exists.

Explanation: An attempt was made to create new key server master keys without a primary key server object. The key server object that is marked as primary is responsible for creating new keys. A primary key server object must exist before key server master keys can be created.

User response: Use the **mkkeyserver** command to designate one key server as the primary key server and run the **chencryption** command again.

CMMVC9138E The command failed because no key servers exist.

Explanation: An attempt was made to create new key server master keys when no key server was known to the system. At least one key server object must exist before key server master keys can be created.

User response: Use the **mkkeyserver** command to create a key server object and run the **chencryption** command again.

CMMVC9139E The command failed because not all key servers are online.

Explanation: An attempt was made to create new key server master keys when the status of a key server object shows that it is not online. All key server objects must be online before key server master keys can be created.

User response: Check the status of all key servers and fix any issues. Run the command again when the status of all key servers is online.

CMMVC9140E The key server could not be deleted because it is the only key server of the enabled key server type.

Explanation: An attempt was made to delete the last remaining key server object for a key server type with a status of **enabled_active**. Deleting this key server would result in loss of access to encryption keys and might cause encrypted objects to go offline.

User response: Use the **chencryption** command to disable key server encryption.

CMMVC9141E The validate command is not valid for key servers. Use the testkeyserver command to validate a key server.

Explanation: An attempt was made to use the **chencryption -keyserver validate** command for key servers, which is not permitted. Only a single key server can be tested at once.

User response: Use the **testkeyserver** command to test individual key server objects.

CMMVC9142E The command failed because another encryption function is already enabled.

Explanation: An attempt was made to enable a second encryption function when another function is already enabled on the system. For example, key server encryption might have been requested when USB encryption was already enabled.

User response: To enable the new encryption function, first use the **chencryption** command to disable the function that is currently enabled.

CMMVC9145E Unable to perform LBA lookup due to unreadable metadata.

Explanation: A reverse lookup failed due to media errors.

User response: Use your host application to validate the data on the volume. After you have located the hard disk errors, restore the missing data from backup.

If the volume is mirrored, and the affected volume is one of the copies, you can use the **repairvdiskcopy**

command with the **-validate** flag to compare the faulty copy to the good one. Use the **-resync** option to restore the data from the good copy to the faulty one.

CMMVC9157E The command has failed because the specified secondary expansion module is offline.

Explanation: The specified secondary expansion module is offline.

User response: This error starts a directed maintenance procedure (DMP) that gives instructions for reseating or replacing the module. If the DMP is not displayed, contact IBM Support.

CMMVC9158E The command has failed because the specified display panel is offline.

Explanation: The specified display panel is offline.

User response: This error starts a directed maintenance procedure (DMP) that gives instructions for replacing the panel. If the DMP is not displayed, contact IBM Support.

CMMVC9159E Host is invalid or does not exist

Explanation: The ID or name of the host that is to be throttled was not valid or was not found.

User response: Use the **lshost** command to create a list of valid hosts. Retry the command with a valid host.

CMMVC9160E Throttle is already associated with this host.

Explanation: The host that was specified in the **mkthrottle** command already has an associated throttle.

User response: Make sure that you specified the correct host. If necessary, use the **lshost** command to create a list of valid hosts and retry the command with a valid host. If you want to change the throttle parameters for this host, use the **chthrottle** command.

CMMVC9161E Host already has an associated host cluster throttle.

Explanation: A host cannot have a throttle if its parent host cluster defines a host cluster throttle.

User response: Make sure that you specified the correct host. If necessary, use the **lshost** command to create a list of valid hosts and retry the command with a valid host. You can tune the system by changing the host cluster throttle to accommodate the requirements of the member hosts by using the **chthrottle** command.

CMMVC9162E Hostcluster is invalid or does not exist.

Explanation: The ID or name of the host cluster that is to be throttled was not valid or was not found.

User response: Use the **lshostcluster** command to create a list of valid host clusters, and then retry the command with a valid host cluster.

CMMVC9163E Throttle is already associated with this hostcluster.

Explanation: The host cluster that was specified in the **mkthrottle** command already has an associated throttle.

User response: Make sure that you specified the correct host cluster. If necessary, use the **lshostcluster** command to create a list of valid host clusters and retry the command with a valid host cluster. If you want to change the throttle parameters for this host cluster, use the **chthrottle** command.

CMMVC9164E Host cluster has a member host which has a throttle defined for it.

Explanation: An attempt was made to define a throttle for a host cluster where throttles are already defined for one or more member hosts. This action is not permitted.

User response: Make sure that you specified the correct host cluster. If so, use the **lsthrottle** command to review the hosts that have throttles. If the information shown is correct and you do not need to define a throttle for the host cluster, no further action is required. If you want to define a throttle for the host cluster, you must remove the individual host throttles by using the **rmthrottle** command. When no throttles are defined for any of the member hosts, retry the command.

CMMVC9165E Mdiskgroup is invalid or does not exist.

Explanation: The ID or name of the storage pool that is to be throttled was not valid or was not found.

User response: Use the **lsmdiskgrp** command to create a list of valid storage pools. Retry the command with a valid storage pool.

CMMVC9166E Throttle is already associated with this mdiskgroup.

Explanation: The storage pool that was specified in the **mkthrottle** command already has an associated throttle.

User response: Make sure that you specified the correct storage pool. If necessary, use the **lsmdiskgrp** command to create a list of valid storage pools and

retry the command with a valid storage pool. If you want to change the throttle parameters for this storage pool, use the **chthrottle** command.

CMMVC9167E Host and hostcluster already have associated throttles.

Explanation: An attempt was made to add a host to a host cluster when throttles were already defined for both the host and the host cluster. Throttles can be defined for the host or for the host cluster, but not for both.

User response: Use the **rmthrottle** command to remove either the host throttle or the host cluster throttle, then retry the **addhostclustermember** command.

CMMVC9168E Seeding host has a throttle associated with it.

Explanation: An attempt was made to create a host cluster when a throttle was defined for one or more of the seeding hosts.

User response: Use the **rmthrottle** command to remove throttles from any seeding hosts, then retry the **mkhostcluster** command.

CMMVC9191E The command failed because the volume is enabled for cloud snapshots.

Explanation: An attempt was made to migrate a volume between storage pools or to add a copy of a volume to different storage pool when the cloud snapshot feature was enabled on the volume. These actions are not permitted.

User response: Make sure that you specified the correct volume. If so, disable the cloud snapshot feature by using the **chvdisk** command and retry the original command.

CMMVC9193E The command failed because the system requested for import does not have data in the account.

Explanation: An attempt was made to configure a cloud account to import data from another system, but the other system does not have data in its account.

User response: Check the list of systems with data in the account by using the **lscloudaccountimportcandidate** command and retry the original command.

CMMVC9194E The command failed because at least one volume is using the account.

Explanation: An attempt was made to change the mode of a cloud account when at least one of the volumes in the system is configured to use the account.

For example, a volume might be configured to use cloud snapshots. The account mode cannot be changed until these volumes are reconfigured to not use the account.

User response: Reconfigure the volumes that are using the account and retry the command.

CMMVC9195E The account cannot be deleted because it is being used by volumes.

Explanation: An attempt was made to delete a cloud account, but volumes on the system are configured to use that account. For example, a volume might be configured to use the cloud snapshot function.

User response: Make sure that you specified the correct account. If so, verify that you want to disconnect the volumes that use the account. If so, disconnect the volumes and retry the command.

CMMVC9196E All available quorum disks are dependent on the sem that you have specified.

Explanation: All quorum disks depend on the specified secondary expander module (SEM). You must configure the system so that at least one of the drives that are allocated to hold quorum remains online when the canister goes offline.

User response: Assign one or more of the drives in the control enclosure as a quorum drive by using the **chquorum** command. After you configure the quorum drives, test for dependencies by using the **lsdependentvdisks** command with the **-sem** option.

CMMVC9197E Cloud snapshot and restore operations are not allowed because of a generation delete abandon. Use rmvolumebackupgeneration to retry the delete operation.

Explanation: A previous generation delete operation was abandoned. You must complete this operation before further snapshot or restore operations can be run.

User response: Rerun the abandoned generation delete operation for the volume. When the operation is complete, new snapshot and restore operations are possible. Use the management GUI or the **lsvolumebackupgeneration** and **rmvolumebackupgeneration** commands to list and remove snapshot generations.

CMMVC9198E A previous rmvolumebackupgeneration -all command was abandoned. Use rmvolumebackupgeneration -all to retry the delete operation.

Explanation: A previous generation delete operation

that used the **-all** option was abandoned. The operation must complete before further snapshots or restores are allowed.

User response: Rerun the abandoned generation delete operation for the volume and include the **-all** option. When the operation is complete, new snapshot and restore operations are possible. Use the management GUI or the **rmvolumebackupgeneration -all** command to remove all snapshot generations for a volume.

CMMVC9199E A previous rmvolumebackupgeneration or cloud restore with -deletelatergenerations command was abandoned. Use rmvolumebackupgeneration specifying the lowest deleting generation to retry the delete operation.

Explanation: One of the following operations was abandoned:

- A generation restore operation with the **-deletelatergenerations** option
- A **rmvolumebackupgeneration** operation with the **-generation** option

The operation must complete before further snapshots or restores are allowed.

User response: Run the **lsvolumebackupgeneration** command and look for one or more generations that are labeled **deleting**. Run the **rmvolumebackupgeneration** operation and specify those generations. When the operation is complete, new snapshot and restore operations are possible.

CMMVC9201E Task failed because volume has a copy that is fully allocated and is part of a Metro Mirror or Global Mirror relationship.

Explanation: Volumes that are part of a Metro Mirror or Global Mirror relationship cannot be expanded or shrunk if they have a copy that is fully allocated.

User response: Either convert all the copies to thin provisioned and wait for the conversion process to complete, or delete the Metro Mirror or Global Mirror relationship, then retry the task.

CMMVC9202E The task failed because the volume is part of a Metro Mirror or Global Mirror relationship to a system which is running software which does not support resize of volumes in a relationship.

Explanation: You can change the size of the volumes that are part of Metro Mirror or Global Mirror relationships only if both systems involved in that

relationship are running software that supports the resizing function.

User response: Either upgrade the remote system or remove the Metro Mirror or Global Mirror relationship, then retry the command.

CMMVC9203E The expandvdiskspace task failed because there is not enough memory available for this feature.

Explanation: The volume is part of Metro Mirror or Global Mirror relationship and requires more memory to accommodate the change recording map.

User response: Increase the memory space available for Remote Copy in the caching I/O group of the volume you want to expand, then retry the task.

CMMVC9204E The task failed because the volume is part of a relationship that is configured to operate in Global Mirror in cycling mode.

Explanation: An attempt was made to change the size of a volume that is part of a Global Mirror relationship while that relationship is configured to operate in multiple cycling mode. This change is not permitted.

User response: Complete one of the following procedures:

- Delete the Global Mirror relationship.
- Remove the relationship from any consistency group, convert the single relationship to Global Mirror non-cycling mode, start the relationship, and wait until the relationship achieves a consistent_synchronized state.

Retry the volume size change.

CMMVC9205E The task failed because the volume is part of a HyperSwap relationship.

Explanation: Volumes cannot be resized if they are part of a HyperSwap relationship.

User response: Make sure that you specified the correct volume. If so, convert the volume to a basic volume by removing a copy at one site, then retry the command.

CMMVC9206E The task failed because the volume is part of a relationship that is not consistent_synchronized.

Explanation: An attempt was made to resize a volume that is part of a remote copy relationship that is not synchronized. A volume that is part of a Metro Mirror or Global Mirror relationship can be resized only when the relationship is synchronized.

User response: Start the relationship and wait for it to

synchronize, and then retry the resize command.

CMMVC9207E The task failed because the volume being expanded is a primary in a Metro Mirror or Global Mirror relationship and would make the primary volume size different than the secondary volume size.

Explanation: The volume that is being expanded is a primary in a Metro Mirror or Global Mirror relationship, and the request to expand is making the volume either larger or smaller than the associated secondary. Volumes in relationships can be expanded by first expanding the secondary by the required amount, and then expanding the primary volume to the same size.

User response: Make sure that the secondary is expanded first, and then adjust the requested size so that the primary volume becomes the same size as the associated secondary volume.

CMMVC9208E The task failed because the volume being expanded is a secondary in a Metro Mirror or Global Mirror relationship and has already been expanded so is larger than the associated primary volume.

Explanation: An attempt was made to expand a secondary volume that was already larger than its associated primary. You cannot expand the secondary volume unless it is the same size as the primary volume.

User response: Either expand the primary volume to the same size as the secondary, or shrink the secondary volume to the same size as the primary. After both volumes are the same size, you can retry the expansion of the secondary volume.

CMMVC9209E The task failed because the volume being shrunk is a secondary in a Metro Mirror or Global Mirror relationship and would make the secondary volume size different than the associated primary volume.

Explanation: The volume that is being shrunk is a secondary in a Metro Mirror or Global Mirror relationship, and the request to shrink would make the volume either larger or smaller than the associated primary volume. Volumes in relationships can be shrunk by first shrinking the primary by the required amount, and then shrinking the secondary volume to the same size.

User response: If you want to shrink the secondary volume, the secondary volume must be larger than the primary volume, and you can shrink it to the size of the primary volume only. When both volumes are the

same size, you can shrink the primary volume first, then shrink the secondary volume to the same size.

CMMVC9210E The task failed because the volume being shrunk is a primary in a Metro Mirror or Global Mirror relationship and has already been shrunk so is already smaller than the associated secondary volume.

Explanation: The volume that is being shrunk is a primary volume that is already smaller than the associated secondary. Volumes in relationships can be shrunk by first shrinking the primary by the required amount, and then shrinking the secondary volume to the same size.

User response: Either shrink the secondary volume to be the same size as the primary volume, or expand the primary volume to be same size as the secondary volume. When both volumes are the same size, you can retry shrinking the primary volume.

CMMVC9211E Cannot perform task because primary is different size to the secondary.

Explanation: The specified command can be completed only on a relationship where the primary and secondary volumes are the same size.

User response: Make sure that you specified the correct relationship. If so, either expand the primary to the same size as the secondary, or shrink the secondary to the same size as the primary. You can then retry the command.

CMMVC9212E Cannot perform task because one or more of the primaries is a different size to the corresponding secondary

Explanation: The specified command can be completed only on a consistency group where all of the relationships consist of primary and secondary volumes that are the same size.

User response: Ensure that all primary volumes are the same size as their corresponding secondary volumes, then retry the command.

Appendix. Accessibility features for SAN Volume Controller

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

These are the major accessibility features for the SAN Volume Controller:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. HTML documents have been tested using JAWS version 15.0.
- This product uses standard Windows navigation keys.
- Interfaces are commonly used by screen readers.
- Keys are discernible by touch, but do not activate just by touching them.
- Industry-standard devices, ports, and connectors.
- You can attach alternative input and output devices.

The SAN Volume Controller online documentation and its related publications are accessibility-enabled. The accessibility features of the online documentation are described in Viewing information in the information center .

Keyboard navigation

You can use keys or key combinations to perform operations and initiate menu actions that can also be done through mouse actions. You can navigate the SAN Volume Controller online documentation from the keyboard by using the shortcut keys for your browser or screen-reader software. See your browser or screen-reader software Help for a list of shortcut keys that it supports.

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

*Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan*

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

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 Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel and Intel Xeon 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 or registered trademarks of Oracle and/or its affiliates.

Linux and the Linux logo is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, 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 product and service names might be trademarks of IBM or other companies.

Index

Special characters

-filtervalue argument xxv

A

about this document
 sending comments xiii
accessibility xi, 945
 repeat rate
 up and down buttons 945
accessing
 publications 945
activatefeature command 419
addcontrolenclosure command 377
addhostclustermember command
 host commands
 addhostclustermember 455
addhostiogrp command 456
addhostport command 456
adding
 nodes 17, 51
addiscsistorageport command 156
addmdisk command 497
addnode command 153
addvdiskaccess command 661
addvdiskcopy command 655
addvolume copy command 653
analyzevdisk command 663
analyzevdiskbysystem command 663
applydrivesoftware command 339
applymdisksoftware command 498
applysoftware command 297, 601
array commands
 charray 89
 charraymember 91
 lsarray 93
 lsarrayinitprogress 100
 lsarraylba 101
 lsarraymember 103
 lsarraymembergoals 106
 lsarraymemberprogress 109
 lsarrayrecommendation 111
 lsarraysyncprogress 114
 lspotentialarraysize 115
 mkarray 117
 mkdistributedarray 120
 overview 89
 recoverarray 123
 recoverarraybysystem 123
 rarray 124
audit log commands
 catauditlog 125
 dumpauditlog 126
 lsauditlogdumps 129
 overview 125
authentication
 SSH logins 1

B

backup and restore commands 129
backup command 130
 help 129
backup commands
 backup 130
 clear 131
 cron 132
backupvolume
 volume commands 664
backupvolume group
 volume commands 665

C

cancellivedump command 493
catauditlog command 125
catterlog command 300
catterlogbyseqnum command 300
cfgportip command 157
changing
 passwords 82
charray command 89
charraymember command 91
chauthservice command 627
chbanner command 163
chbootdrive command 605
chcloudaccountawss3
 clustered system commands 137
chcloudaccountswift
 clustered system commands 139
chcluster command 165
chcontroller commands 331
chcurrentuser command 630
chdnserver
 Email and event notification
 commands 300
chdrive command 342
chemail command 359
chemailserver command 361
chemailuser command 362
chenclosure command 378
chenclosurecanister command 378
chenclosuredisplaypanel
 clustered system commands 380
chenclosurepsu command 380
chenclosuresem
 clustered system commands 381
chenclosureslot command 381
chencryption command 413
cherrstate command 300
cheventlog command 301
chfcconsistgrp command 429
chfcmap command 429
chhost command 458
chhostcluster command
 host commands
 chhostcluster 459
chiogrp command 165

chkeyserver
 Key server commands 485
chkeyserverisklm command
 clustered system commands
 chkeyserverisklm 486
chldap command 630
chldapserver command 633
chlicense command 420
chmdisk command 498
chmdiskgrp command 517
chnode / chnodecanister command 168
chnodebattery command 171
chnodebootdrive command 172
chnodehw / chnodecanisterhw
 command 173
chnodeled command 605
chpartnership command 535
chquorum command 174
chrconsistgrp command 538
chrrelationship command 540
chsecurity command 176
chserviceip command 607
chsite command 176
chsnmpserver command 363
chsyslogserver command 301
chsystem command 177
chsystemcert command 184
chsystemip command 187
chthrottle command 189
chuser command 635
chusergrp command 636
chvdisk command 666
chvolume group
 volume commands 670
chvpd command 609
chwwnn command 611
clear command 131
 help 129
cleardumps command 190, 601
clearerrlog command 302
CLI (command-line interface)
 configuring PuTTY 3
 getting started 11
 preparing SSH client on AIX or
 Linux 6
 preparing SSH client on Windows 2
 using to update clustered system
 license 12
CLI commands
 chcurrentuser 78
 chfcmap 36
 chlicense 12
 chsystem
 changing relationship
 bandwidth 72
 changing system gateway
 address 72
 modifying system IP address 71
 chsystemip
 modifying system IP address 71
 chuser 78

CLI commands (continued)

- chusergrp 77
- lscurrentuser 78
- lsfcconsistgrp 36, 37
- lsfcmap 33, 36
- lslicense 12
- lssystem
 - changing relationship
 - bandwidth 72
 - changing system gateway
 - address 72
 - displaying clustered system
 - properties 13
- lssystemip
 - modifying system IP address 71
- lsuser 78
- lsusergrp 77
- lsvdisk 33
- mkfcconsistgrp 36
- mkfcmap 33
- prestartfcconsistgrp 37
- setlocale 83
- startfcconsistgrp 37
- cluster date and time
 - setting 12
- Cluster diagnostic and service-aid
 - commands
 - lsupdate 322
- cluster error log
 - displaying 300
- clustered syetem commands
 - rmportip 288
- clustered system
 - authentication
 - configuring clustered system
 - iSCSI 75
 - configuring for iSCSI 73
 - configuring iSCSI authentication 75
- clustered system commands
 - addiscsistorageport 156
 - addnode 153
 - cfgportip 157
 - chbanner 163
 - chcloudaccountawss3 137
 - chcloudaccountswift 139
 - chenclosuresdisplaypanel 380
 - chenclosuresem 381
 - chencryption 413
 - chiogrp 165
 - chkeyserverisklm 486
 - chnode/ chnodecanister 168
 - chnodebattery 171
 - chsystem 177
 - chsystemcert 184
 - chsystemip 187
 - chthrottle 189
 - cleardumps 190
 - cpdumps 191
 - detectiscsistorageportcandidate 193
 - detectmdisk 500
 - host 587
 - lscloudaccount 142
 - lscloudaccountimportcandidate 146
 - lscloudaccountusage 144
 - lsclosuresdisplaypanel 394
 - lsclosuressem 400
 - lsiscsistorageport 201

clustered system commands (continued)

- lsiscsistorageportcandidate 205
- lskeyserver 487
- lskeyserverisklm 488
- lssystemcert 270
- lstargetportfc 280
- lsthrottle 261
- mkcloudaccountawss3 147
- mkcloudaccountswift 148
- mkcluster 616
- mkthrottle 282
- ping 284
- rmcloudaccount 150
- rmiscsistorageport 285
- rmkeyserver 491
- rmthrottle 289
- satask mkcluster 617
- setpwreset 290
- setsystemtime 289
- settimezone 291
- stopsystem 294
- testcloudaccount 150
- testkeyserver 492
- traceroute 600
- clustered system diagnostic and
 - service-aid commands
 - applysoftware 297
 - cheventlog 301
 - clearerrlog 302
 - dumpperlog 303
 - finderr 304
 - overview 297
 - setlocale 328, 620
 - svqueryclock 329
 - writesernum 329
- clustered systems 275, 630, 633, 637, 638, 643, 648, 649
 - configuring iSCSI alias 74
 - modifying iSCSI alias 74
 - properties 13
 - updating
 - license 12
 - viewing
 - license 12
- clustered systems commands
 - startstats 292
- clustered sytem commands
 - rmnode / rmnodecanister 285
- clusters
 - error logs 84
 - logs 83
 - viewing feature logs 83
- command 352, 354
- 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 2
 - using to update clustered system
 - license 12
 - using to view clustered system
 - license 12
- Command-line interface messages
 - overview 775

- commands 630, 633, 637, 638, 643, 648, 649
 - activatefeature 419
 - addcontrolenclosure 377
 - addhostiogrp 456
 - addhostport 456
 - addmdisk 497
 - addnode 153
 - addvdiskaccess 661
 - addvdiskcopy 655
 - applydrivesoftware 339
 - applymdisksoftware 498
 - applysoftware 297, 601
 - backup 130
 - cancellivedump 493
 - catauditlog 125
 - caterlog 300
 - caterlogbyseqnum 300
 - cfgportip 157
 - charray 89
 - charraymember 91
 - chauthservice 627
 - chbootdrive 605
 - chcluster 165
 - chcontroller 331
 - chcurrentuser 630
 - chdrive 342
 - chemail 359
 - chemailserver 361
 - chemailuser 362
 - chenclosure 378
 - chenclosurecanister 378
 - chenclosurepsu 380
 - chenclosureslot 381
 - chencryption 413
 - cherrstate 300
 - cheventlog 301
 - chfcconsistgrp 429
 - chfcmap 429
 - chhost 458
 - chiogrp 165
 - chlicense 420
 - chmdisk 498
 - chmdiskgrp 517
 - chnode/ chnodecanister 168
 - chnodebattery 171
 - chnodebootdrive 172
 - chnodehw / chnodecanisterhw 173
 - chnodeled 605
 - chpartnership 535
 - chquorum 174
 - chrconsistgrp 538
 - chrrelationship 540
 - chsecurity 176
 - chserviceip 607
 - chsite 176
 - chsnmpserver 363
 - chsyslogserver 301
 - chsystem 177
 - chsystemcert 184
 - chsystemip 187
 - chuser 635
 - chusergrp 636
 - chvdisk 666
 - chvpd 609
 - chwnnn 611
 - clear 131

commands (*continued*)

cleardumps 190, 601
 clearerrlog 302
 cpdumps 191
 cpfiles 612
 cron 132
 deactivatefeature 422
 detectmdisk 500
 dumpallmdiskbadblocks 502
 dumpauditlog 126
 dumperrlog 303, 601
 dumpinternallog 613
 dumpmdiskbadblocks 502
 exit 603
 expandvdisksize 671
 finderr 304
 help 194
 includemdisk 504
 installsoftware 613
 leavecluster 614
 licensing 419
 livedump 493
 ls2145dumps 483, 603
 lsarray 93
 lsarrayinitprogress 100
 lsarraylba 101
 lsarraymember 103
 lsarraymembergoals 106
 lsarraymemberprogress 109
 lsarrayrecommendation 111
 lsarraysyncprogress 114
 lsauditlogdumps 129
 lsbootdrive 588
 lscimomdumps 304, 603
 lsclustervpd 603
 lscmdstatus 589
 lscontrolenclosurecandidate 389
 lscontroller 332
 lscontrollerdependentvdisks 336
 lscopystatus 305
 lscurrentuser 636
 lsdependentvdisks 674
 lsdiscoverystatus 195
 lsdrive 343
 lsdriveclass 349
 lsdrive1ba 351
 lsdumps 305
 lsemailserver 365
 lsemailuser 366
 lsenclature 383
 lsenclaturebattery 386
 lsenclaturecanister 390
 lsenclaturechassis 393
 lsenclaturefanmodule 396
 lsenclaturepsu 398
 lsenclatureslot 402
 lsenclaturestats 406
 lsencryption 415
 lserrlogbyfcconsistgrp 308
 lserrlogbyfcmap 308
 lserrlogbyhost 308
 lserrlogbyiogrp 308
 lserrlogbymdisk 308
 lserrlogbymdiskgp 308
 lserrlogbynnode 308
 lserrlogbyrcconsistgrp 309
 lserrlogbyrcrelationship 309

commands (*continued*)

lserrlogbyvdisk 309
 lserrlogdumps 309, 603
 lseventlog 309
 lsfabric 197
 lsfcconsistgrp 431
 lsfcmap 434
 lsfcmapcandidate 438
 lsfcmapdependentmaps 439
 lsfcmapprogress 438
 lsfcportcandidate 200
 lsfeature 423
 lsfeaturedumps 309, 603
 lsfiles 590
 lsfreeextents 519
 lshardware 591
 lshost 461
 lshostiogrp 470
 lshostvdiskmap 675
 lsiogrp 206
 lsiogrpcandidate 211
 lsiogrpghost 210
 lsiostatsdumps 211, 603
 lsio tracedumps 211, 603
 lsiscsiauth 471
 lslicense 425
 lslivedump 493, 494
 lsmdisk 504
 lsmdiskcandidate 511
 lsmdiskdumps 510, 603
 lsmdiskextent 513
 lsmdiskgrp 519
 lsmdisklba 510
 lsmdiskmember 514
 lsmetadataavdisk 677
 lsmigrate 581
 lsnodebattery 216
 lsnodecandidate 219
 lsnode dependentvdisks 221
 lsnodehw / lsnodecanisterhw 221
 lsnodestats / lsnodecanisterstats 223
 lsnodevpd / lsnodecanistervpd 231
 lspartnership command 544
 lspartnershipcandidate 548
 lspportfc 250
 lspportip 242
 lspportsas 252
 lspportusb 240
 lspotentialarraysize 115
 lsquorum 254
 lsrconsistgrp 549
 lsrelationship 552
 lsrelationshipcandidate 557
 lsrelationshipprogress 558
 lsrepairsevdiskcopyprogress 678
 lsrepairdiskcopyprogress 680
 lsrmvdiskdependentmaps 440
 lsroute 256
 lssasportcandidate 258
 lssecurity 259
 lsservicenodes 593
 lsservicerecommendation 594
 lsservicestatus 315, 595
 lssevdiskcopy 681
 lssite 260
 lssnmpserver 366
 lssoftwaredumps 321, 605

commands (*continued*)

lssyslogserver 320
 lssystem 263
 lssystemcert 270
 lssystemip 272
 lssystemstats 275
 lstimezones 257
 lsupdate 322
 lsuser 640
 lsusergp 641
 lsvdisk 687
 lsvdiskaccess 702
 lsvdiskcopy 707
 lsvdiskdependentmaps 713
 lsvdiskextent 714
 lsvdiskfcmapcopies 715
 lsvdiskfcmappings 716
 lsvdiskhostmap 717
 lsvdisklba 718
 lsvdiskmember 720
 lsvdiskprogress 721
 lsvdisksyncprogress 722
 metadata 615
 migrateexts 582
 migrate toimage 583
 migratevdisk 585
 mkarray 117
 mkcluster 616
 mkdistributedarray 120
 mkemailserver 367
 mkemailuser 368
 mkfcconsistgrp 441
 mkfcmap 442
 mkfcpartnership 558
 mkhost 472
 mkippartnership 559
 mkmdiskgrp 527
 mkmetadavdisk 733
 mkquorumapp 282
 mkrcconsistgrp 561
 mkrcrelationship 562
 mksnmpserver 370
 mksyslogserver 325
 mkuser 644
 mkusergp 645
 mkvdisk 734
 mkvdiskhostmap 743
 movevdisk 751
 overridequorum 617
 ping 284
 prestartfcconsistgrp 444, 449
 prestartfcmap 446
 recover 132
 recoverarray 123
 recoverarraybysystem 123
 recovervdisk 753
 recovervdiskbyiogrp 754
 recovervdiskbysystem 754
 repairsevdiskcopy 755
 repairvdiskcopy 756
 rescuene 618
 resetleds 411
 resetpassword 618
 restartservice 619
 restore 133, 135
 rmarray 124
 rmemailserver 371

- commands (*continued*)
 - rmemailuser 371
 - rmfconsistgrp 447
 - rmfcmap 448
 - rmhost 476
 - rmhostiogr 480
 - rmhostport 481
 - rmmdisk 531
 - rmmdiskgrp 532
 - rmmetadatavdisk 761
 - rmnode / rmnodecanister 285
 - rmpartnership 566
 - rmportip 288
 - rmrconsistgrp 567
 - rmrrelationship 567
 - rmsnmpserver 372
 - rmsyslogserver 327
 - rmuser 648
 - rmusergrp 649
 - rmvdisk 759
 - rmvdiskaccess 763
 - rmvdiskcopy 762
 - rmvdiskhostmap 763
 - sataisk mkcluster 617
 - sendinventoryemail 372
 - service information 587
 - service task 605
 - setlocale 328, 620
 - setpacedccu 621
 - setpwdreset 290
 - setquorum 516
 - setsystemtime 289
 - settempsshkey 621
 - settimezone 291
 - showtimezone 291
 - shrinkvdisksize 766
 - snap 622
 - splitvdiskcopy 772
 - startemail 373
 - startfcconsistgrp 449
 - startfcmap 450
 - starttrconsistgrp 568
 - starttrrelationship 571
 - startservice 623
 - startstats 292
 - stopcluster 293
 - stopemail 374
 - stopfcconsistgrp 452
 - stopfcmap 453
 - stopnode 623
 - stopprconsistgrp 574
 - stopprrelationship 576
 - stopservice 624
 - stopsystem 294
 - svconfig 129
 - svqueryclock 329
 - switchrconsistgrp 578
 - switchrrelationship 579
 - t3recovery 625
 - testemail 374
 - triggerenclosedump 412
 - triggerlivedump 495
 - triggermddiskdump 517
 - user management 627
 - writesernum 329
- commands/ lsnodecanister
 - lsnode 211

- comments, sending xiii
- communications
 - determining between hosts and volumes 49
- configuring
 - iSNS server address 75
 - PuTTY 3
 - remote authentication service using CLI 76
 - remote authentication service with Lightweight Directory Access Protocol (LDAP) using CLI 76
- Connecting to the CLI using OpenSSH 8
- consistency group
 - deleting FlashCopy 39
 - stoppingFlashCopy 39
- consistency groups, active-active
 - creating 43
 - deleting 45
 - modifying 43
 - starting and stopping 44
- consistency groups, Global Mirror
 - creating 43
 - deleting 45
 - modifying 43
 - starting and stopping 44
- consistency groups, Metro Mirror
 - creating 43
 - deleting 45
 - modifying 43
 - starting and stopping 44
- controller commands
 - chcontroller 331
 - overview 331
- controllers
 - changing 331
 - command 331, 332
- cpdumps command 191
- cpfiles command 612
- creating
 - host mappings 33
- creating users 8
- cron command 132
- help 129
- current time zone 291

D

- data migration progress
 - viewing 581
- date and time
 - setting cluster 12
- deactivatefeature command 422
- deleting
 - nodes 68
- dependent maps
 - viewing 439
- detectiscsistorageportcandidate command 193
- detectmdisk command 500
- determining
 - communications between hosts and volumes 49
- diagnostic and service-aid commands
 - clearerrlog
 - clustered system 302

- diagnostic and service-aid commands (*continued*)
 - cluster
 - svqueryclock 329
 - clustered system 297
 - applysoftware 297
 - cheventlog 301
 - setlocale 328, 620
 - writesernum 329
 - dumperrlog
 - cluster 303
 - finderr
 - clustered system 304
 - overview 297
- discovering
 - managed disks 20
- disks
 - migrating 64
 - migrating image mode 68
- documentation
 - improvement xiii
- drive commands
 - applydrivesoftware 339
 - chdrive 342
 - lsdrive 343
 - lsdriveclass 349
 - lsdrivelba 351
 - overview 339
- dump files
 - listing 483, 603
 - lsfeaturedumps 309
- dumpallmdiskbadblocks command 502
- dumpauditlog command 126
- dumperrlog command 303, 601
- dumpinternallog command 613
- dumpmdiskbadblocks command 502

E

- email
 - inventory reports 81
 - setting up event notification 81
- email and event notification commands
 - chemailserver 361
 - chsnmpserver 363
 - chsyslogserver 301
 - mkemailserver 367
 - mksnmpserver 370
 - mksyslogserver 325
 - rmemailserver 371
 - rmsnmpserver 372
 - rmsyslogserver 327
- Email and event notification commands
 - chdnsserver 300
 - lsdnsserver 307
 - mkdnsserver 325
 - rmdnsserver 327
- email commands
 - chemail 359
 - chemailuser 362
 - lsemailuser 366
 - mkemailuser 368
 - overview 359
 - rmemailuser 371
 - sendinventoryemail 372
 - startemail 373
 - stopemail 374

- email commands (*continued*)
 - testemail 374
- email servers
 - setting up
 - CLI 82
- enclosure commands
 - chenclosure 378
 - lsenclosure 383
 - lsenclosurecanister 390
 - lsenclosurepsu 398
 - lsenclosureslot 402
 - overridequorum 617
 - overview 377
- error log dump files
 - viewing 603
- error notification
 - SYSLOG 80
- event notification commands
 - overview 359
- exit command 603
- expanding
 - volume 60
- expandvdisize command 671
- extent allocation
 - viewing 513
- extents
 - migrating
 - using the CLI (command-line interface) 63

F

- featurization settings 425
- filtering
 - FlashCopy
 - consistency groups 431
 - mappings 434, 440, 713, 715
- finderr command 304
- FlashCopy
 - consistency group
 - deleting using CLI 39
 - stopping using CLI 39
 - consistency groups
 - creating using CLI 36
 - preparing using the CLI 37
 - starting using the CLI 37
 - deleting consistency group 39
 - deleting mapping 36
 - mapping
 - deleting using CLI 36
 - stopping 35
 - mappings
 - adding to consistency group 36
 - creating using CLI 33
 - memory 26
 - stopping consistency group 39
- FlashCopy commands
 - chfcconsistgrp 429
 - chfcmap 429
 - mkfcconsistgrp 441
 - mkfcmap 442
 - overview 429
 - prestartfcconsistgrp 444, 449
 - prestartfcmap 446
 - rmfcconsistgrp 447
 - rmfcmap 448
 - startfcconsistgrp 449

- FlashCopy commands (*continued*)
 - startfcmap 450
 - stopfcconsistgrp 452
 - stopfcmap 453
- FlashCopy progress 438
- free extents 519
- front panel
 - password 14

G

- gateway address
 - changing 72
- Generating an SSH key pair using OpenSSH 7
- getting started
 - using the CLI (command-line interface) 11
 - using the command-line interface (CLI) 11
- Global Mirror
 - memory 26
- Global Mirror commands
 - chpartnership 535
 - chrcconsistgrp 538
 - chrcrelationship 540
 - mkrcconsistgrp 561
 - mkrcrelationship 562
 - overview 535
 - rmpartnership 566
 - rmrcconsistgrp 567
 - rmrcrelationship 567
 - startrrconsistgrp 568
 - startrrrelationship 571
 - stoprrconsistgrp 574
 - stoprrrelationship 576
 - switchrrconsistgrp 578
 - switchrrrelationship 579

H

- help command 194
- host
 - clustered system commands 587
- host commands
 - addhostclustermember 455
 - addhostiogrp 456
 - addhostport 456
 - addvdiskaccess 661
 - chhost 458
 - chhostcluster 459
 - lshostcluster 465
 - lshostclustermember 467
 - lshostclustervolumemap 468
 - mkhost 472
 - mkhostcluster 474
 - mkvolumehostclustermember 475
 - movevdisk 751
 - overview 455
 - rmhost 476
 - rmhostcluster 477
 - rmhostclustermember 478
 - rmhostiogrp 480
 - rmhostport 481
 - rmvdiskaccess 763
 - rmvolumehostclustermember 479

- host I/O group 470
- host objects
 - configuring using CLI 31
- hosts
 - commands 455
 - determining volume names 49
 - mapping volumes 33
 - viewing 461
- HyperSwap commands
 - overview 535

I

- image-mode volumes
 - converting to managed mode
 - using CLI (command-line interface) 67
- includemdisk command 504
- information commands 440, 630, 633, 637, 638, 643, 648, 649, 715
 - addcontrolenclosure 377
 - caterlog 300
 - caterlogbyseqnum 300
 - chbootdrive 605
 - chenclosurecanister 378
 - chenclosurepsu 380
 - chenclosureslot 381
 - chnodebootdrive 172
 - chnodehw / chnodecanisterhw 173
 - chsecurity 176
 - chsite 176
 - chvpd 609
 - ls2145dumps 483
 - lsbootdrive 588
 - lscimomdumps 304
 - lscontrolenclosurecandidate 389
 - lscontroller 332
 - lscopystatus 305
 - lscurrentuser 636
 - lsdependentvdisks 674
 - lsdiscoverystatus 195
 - lsdumps 305
 - lsemailer 365
 - lsenclosurebattery 386
 - lsenclosurechassis 393
 - lsenclosurefanmodule 396
 - lsenclosurestats 406
 - lscryption 415
 - lserrlogbyfcconsistgrp 308
 - lserrlogbyfcmap 308
 - lserrlogbyhost 308
 - lserrlogbyiogrp 308
 - lserrlogbymdisk 308
 - lserrlogbymdiskgp 308
 - lserrlogbynode 308
 - lserrlogbyrconsistgrp 309
 - lserrlogbyrrelationship 309
 - lserrlogbyvdisk 309
 - lserrlogdumps 309
 - lseventlog 309
 - lsfabric 197
 - lsfcconsistgrp 431
 - lsfcmap 434
 - lsfcmapcandidate 438
 - lsfcmapdependentmaps 439
 - lsfcmapprogress 438
 - lsfeaturedumps 309

information commands (continued)

- lsfreeextents 519
- lshardware 591
- lshost 461
- lshostiogrp 470
- lshostvdiskmap 675
- lsiogrp 206
- lsiogrpcandidate 211
- lsiogrphost 210
- lsiostatsdumps 211
- lsiotracedumps 211
- lsiscsiauth 471
- lslicense 425
- lsmdisk 504
- lsmdiskcandidate 511
- lsmdiskdumps 510
- lsmdiskextent 513
- lsmdiskgrp 519
- lsmdisklba 510
- lsmdiskmember 514
- lsmigrate 581
- lsnodebattery 216
- lsnodecandidate 219
- lsnodependentvdisks 221
- lsnodehw / lsnodecanister 221
- lsnodestats / lsnodecanisterstats 223
- lsnodevpd / lsnodecanistervpd 231
- lspartnership command 544
- lspartnershipcandidate 548
- lsportfc 250
- lsportip 242
- lsportsas 252
- lsportusb 240
- lsquorum 254
- lsrconsistgrp 549
- lsrrelationship 552
- lsrrelationshipcandidate 557
- lsrrelationshipprogress 558
- lsroute 256
- lssecurity 259
- lssite 260
- lssnmpserver 366
- lssoftware dumps 321
- lssyslogserver 320
- lssystem 263
- lssystemip 272
- lssystemstats 275
- lstimezones 257
- lsuser 640
- lsusergrp 641
- lsvdisk 687
- lsvdiskaccess 702
- lsvdiskdependentmaps 713
- lsvdiskextent 714
- lsvdiskfcmappings 716
- lsvdiskhostmap 717
- lsvdisklba 718
- lsvdiskmember 720
- lsvdiskprogress 721
- mkquorumapp 282
- overview 483
- resetleds 411
- showtimezone 291
- stopcluster 293
- triggerenclosedump 412

information commands/ lsnodecanister

- lsnode 211

- installsoftware command 613
- inventory commands
 - chemail 359
 - chsystem 177
 - mkemailuser 368
 - rmemailuser 371
 - sendinventoryemail 372
 - startemail 373
 - stopemail 374
 - testemail 374
- IP addresses
 - changing 71
- iSCSI alias
 - configuring 74
 - modifying 74
- iSNS server address
 - configuring 75

K

- Key server commands
 - chkeyserver 485
 - mkkeyserver 491
- keyboards
 - accessibility features 945
- Knowledge Center xi

L

- language
 - changing locale 83
- ldapserver command 638
- license
 - changing settings 420
 - updating
 - using the CLI (command-line interface) 12
 - viewing 425
- licensing commands 419
 - chlicense 420
 - dumpinternallog 613
- Lightweight Directory Access Protocol (LDAP) using CLI
 - configuring remote authentication service 76
- list dump command 15
- livedump commands 493
 - cancellivedump 493
 - lslivedump 493, 494
 - triggerlivedump 495
- locale
 - changing 83
- ls2145dumps command 483, 603
- lsarray command 93
- lsarrayinitprogress command 100
- lsarraylba command 101
- lsarraymember command 103
- lsarraymembergoals command 106
- lsarraymemberprogress command 109
- lsarrayrecommendation command 111
- lsarraysyncprogress command 114
- lsauditlogdumps command 129
- lsbootdrive command 588
- lscimomdumps command 304, 603
- lscloudaccount
 - clustered system commands 142

- lscloudaccountimportcandidate
 - clustered system commands 146
- lscloudaccountusage
 - clustered system commands 144
- lsclustervpd command 603
- lscmdstatus command 589
- lscontrolenclosurecandidate
 - command 389
- lscontroller command 332
- lscontrollerdependentvdisks
 - command 336
- lscopystatus command 305
- lscurrentuser command 636
- lsdependentvdisks command 674
- lsdiscoverystatus command 195
- lsdnsserver
 - Email and event notification
 - commands 307
- lsdrive command 343
- lsdriveclass command 349
- lsdrivelba command 351
- lsdriveprogress 352
- lsdriveprogress command 352
- lsdriveupgradepress 354
- lsdriveupgradepress command 354
- lsdumps command 305
- lsemailer command 365
- lsemailer command 366
- lsclosure command 383
- lsclosurebattery command 386
- lsclosurecanister command 390
- lsclosurechassis command 393
- lsclosuredisplaypanel
 - clustered system commands 394
- lsclosurefanmodule command 396
- lsclosurepsu command 398
- lsclosuresem
 - clustered system commands 400
- lsclosureslot command 402
- lsclosurestats command 406
- lscryption command 415
- lserrlogbyfconsistgrp command 308
- lserrlogbyfcmmap command 308
- lserrlogbyhost command 308
- lserrlogbyiogrp command 308
- lserrlogbymdisk command 308
- lserrlogbymdiskgp command 308
- lserrlogbyn timer command 308
- lserrlogbyrconsistgrp command 309
- lserrlogbyrrelationship command 309
- lserrlogbyvdisk command 309
- lserrlogdumps command 309, 603
- lseventlog command 309
- lsfabric command 197
- lsfconsistgrp command 431
- lsfcmmap command 434
- lsfcmmapcandidate command 438
- lsfcmmapdependentmaps command 439
- lsfcmmapprogress command 438
- lsfcpportcandidate command 200
- lsfeature command 423
- lsfeaturedumps command 309
- lsfeaturedumps commands 603
- lsfiles command 590
- lsfreeextents command 519
- lshardware command 591
- lshost command 461

- lshostcluster command
 - host commands
 - lshostcluster 465
- lshostclustermember command
 - host commands
 - lshostclustermember 467
- lshostclustervolumemap command
 - host commands
 - lshostclustervolumemap 468
- lshostigrp command 470
- lshostvdiskmap command 675
- lsiogrp command 206
- lsiogrpcandidate command 211
- lsiogrpghost command 210
- lsiostatsdumps command 211, 603
- lsiotracedumps command 211
- lsiotracedumps commands 603
- lsisciauth command 471
- lsiscsistorageport command 201
- lsiscsistorageportcandidate command 205
- lskeyserver
 - clustered system commands 487
- lskeyserverisklm command
 - clustered system commands
 - lskeyserverisklm 488
- lsldap command 637
- lslicense command 425
- lslivedump command 493, 494
- lsmdisk command 504
- lsmdiskcandidate command 511
- lsmdiskdumps command 510, 603
- lsmdiskextent command 513
- lsmdiskgrp command 519
- lsmdisklba command 510
- lsmdiskmember command 514
- lsmetadavdisk command 677
- lsmigrate command 581
- lsnode command 211
- lsnodebattery command 216
- lsnodecandidate command 219
- lsnodecandidatevdisks command 221
- lsnodehw / lsnodecanisterhw command 221
- lsnodestats / lsnodecanisterstats command 223
- lsnodevpd / lsnodecanistervpd command 231
- lspartnership command 544
- lspartnershipcandidate command 548
- lsportfc command 250
- lsportip command 242
- lsportsas command 252
- lsportusb command 240
- lspotentialarraysize command 115
- lsquorum command 254
- lsrcconsistgrp command 549
- lsrcrelationship command 552
- lsrcrelationshipcandidate command 557
- lsrcrelationshipprogress command 558
- lsrepairsevdiskcopyprogress command 678
- lsrepairvdiskcopyprogress command 680
- lsrmvdiskdependentmaps command 440
- lsroute command 256
- lssasportcandidate command 258

- lssecurity command 259
- lsservicenodes command 593
- lsservicerecommendation command 594
- lsservicestatus command 315, 595
- lssevdiskcopy command 681
- lssite command 260
- lssnmpserver command 366
- lssoftwaredumps command 321, 605
- lssyslogserver command 320
- lssystem command 263
- lssystemcert command 270
- lssystemip command 272
- lssystemstats command 275
- lstargetportfc command 280
- lsthrottle command 261
- lstimezones command 257
- lsupdate command 322
- lsuser command 640
- lsusergrp command 641
- lsvdisk command 687
- lsvdiskaccess command 702
- lsvdiskanalysis command 703
- lsvdiskanalysisprogress command 706
- lsvdiskcopy command 707
- lsvdiskdependentmaps command 713
- lsvdiskextent command 714
- lsvdiskfcmappings command 715
- lsvdiskfcmappings command 716
- lsvdiskhostmap command 717
- lsvdisklba command 718
- lsvdiskmember command 720
- lsvdiskprogress command 721
- lsvdisksyncprogress command 722
- lsvolumebackup
 - volume commands 723
- lsvolumebackupgeneration
 - volume commands 725
- lsvolumebackupprogress
 - volume commands 727
- lsvolumebackupgroup
 - volume commands 729
- lsvolumebackupprogress
 - volume commands 731

M

- maintaining
 - passwords 14
- managed disk commands
 - applymdisksoftware 498
 - chmdisk 498
 - chquorum 174
 - includemdisk 504
 - lsquorum 254
 - mkquorumapp 282
 - overview 497
 - setquorum 516
 - triggermdiskdump 517
- managed disks
 - viewing disks 504, 510
 - viewing groups 519
- managed disks (MDisks)
 - adding 24
 - discovering 20
 - rebalancing access 20
 - volume relationships 50

- managed mode volumes
 - converting from image mode
 - using the CLI (command-line interface) 67
- mapping
 - deleting FlashCopy 36
- master console
 - configuration 2
- MDisk commands
 - dumpallmdiskbadblocks 502
 - dumpmdiskbadblocks 502
- MDisks (managed disks)
 - adding 24
 - volume relationships 50
- MDisks *See* managed disks 497, 517
- metadata command 615
- Metro Mirror
 - memory 26
- Metro Mirror and Global Mirror commands
 - mkfcpartnership 558
 - mkippartnership 559
- Metro Mirror commands
 - chpartnership 535
 - chrcconsistgrp 538
 - mkrrcconsistgrp 561
 - mkrrccrelationship 562
 - overview 535
 - rmpartnership 566
 - rmrcconsistgrp 567
 - rmrccrelationship 567
 - startrrcconsistgrp 568
 - startrrccrelationship 571
 - stoprrcconsistgrp 574
 - stoprrccrelationship 576
 - switchrrcconsistgrp 578
 - switchrrccrelationship 579
- migrateexts command 582
- migratetoimage command 583
- migratevdisk command 585
- migratingvolumes
 - extents
 - using the CLI (command-line interface) 63
- migration 581
- migration commands
 - migrateexts 582
 - migratetoimage 583
 - migratevdisk 585
 - overview 581
- mkarray command 117
- mkcloudaccountawss3
 - clustered system commands 147
- mkcloudaccountswift
 - clustered system commands 148
- mkcluster command 616
 - See* sastask mkcluster
- mkdistributedarray command 120
- mkdnsserver
 - Email and event notification commands 325
- mkemailserver command 367
- mkemailuser command 368
- mkfcccconsistgrp command 441
- mkfcccmap command 442
- mkfcccpartnership command 558
- mkhost commands 472

- mkhostcluster command
 - host commands
 - mkhostcluster 474
- mkimagevolume command 749
- mkippartnership command 559
- mkkeyserver command
 - Key server commands 491
- mkldapserver command 643
- mkmdiskgrp command 527
- mkmetadavdisk command 733
- mkquorumapp command 282
- mkrcconsistgrp command 561
- mkrcrelationship command 562
- mksnmpserver command 370
- mksyslogserver command 325
- mkthrottle command 282
- mkuser command 644
- mkusergrp command 645
- mkvdisk commands 734
- mkvdiskhostmap command 743
- mkvolume command 745
- mkvolumegroup
 - volume commands 748
- mkvolumehostclustermap command
 - host commands
 - mkvolumehostclustermap 475
- modifying system IP address
 - chsystemip 71
- movevdisk command 751

N

- navigation
 - accessibility 945
- nodes
 - adding 17, 51, 153
 - addnode command 153
 - changing 168
 - chnode/ chnodecanister
 - command 168
 - deleting 68, 285
 - lsnodestats / lsnodestats
 - command 223
 - removing 68
 - returning to the system 57
 - rmnode / rmnodecanister
 - command 285
 - statistics 223
 - viewing 211
 - general details 20

O

- openssh scp
 - copying software update files 9
- OpenSSH, Connecting to the CLI
 - using 8
- OpenSSH, Generating an SSH key pair
 - using 7
- overridequorum command 617
- overview
 - array commands 89
 - audit log commands 125
 - backup and restore commands 129
 - cloud account commands 137
 - cluster commands 153

- overview (*continued*)
 - clustered system diagnostic and
 - service-aid commands 297
 - controller commands 331
 - drive commands 339
 - dumps commands 15
 - email commands 359
 - enclosure commands 377
 - event notification commands 359
 - FlashCopy commands 429
 - host commands 455
 - information commands 483
 - key server commands 485
 - licensing commands 419
 - managed disk commands 497
 - migration commands 581
 - secure shell 1
 - security commands 413
 - service mode commands 601
 - service mode information
 - commands 603
 - storage pool commands 517
 - user management commands 627

P

- partnerships, Global Mirror
 - creating 45
 - deleting 47
 - modifying 46
 - starting and stopping 47
- partnerships, Metro Mirror
 - creating 45
 - deleting 47
 - modifying 46
 - starting and stopping 47
- passwords
 - changing 82
 - front panel 14
- ping command 284
- plink utility
 - running 4
- port IP addresses
 - configuring 73
- powering off
 - system 84
- prestartfcconsistgrp command 444
- prestartfcmap command 446
- publications
 - accessing 945
- PuTTY
 - configuring 3
 - generating an SSH key pair 2
 - running the plink utility 4
- PuTTY pscp
 - copying software update files 9
- PuTTY session
 - configuring for the CLI 3

Q

- quorum disks
 - setting with CLI 25

R

- RAID
 - memory 26
- reader feedback, sending xiii
- rebalancing
 - managed disks (MDisks) access 20
- recover command 132
 - help 129
- recover commands
 - recover 132
- recoverarray command 123
- recoverarraybysystem command 123
- recovering
 - offline volumes
 - using CLI 56
- recovervdisk command 753
- recovervdiskbyiogrp command 754
- recovervdiskbysystem command 754
- related information xi
- relationships, active-active
 - creating 40
 - deleting 43
 - displaying 42
 - modifying 41
 - starting and stopping 41
- relationships, Global Mirror
 - creating 40
 - deleting 43
 - displaying 42
 - modifying 41
 - starting and stopping 41
 - switching 42
- relationships, Metro Mirror
 - creating 40
 - deleting 43
 - displaying 42
 - modifying 41
 - starting and stopping 41
 - switching 42
- remote authentication
 - configuring using CLI 76
- removing
 - nodes 68
- repairing
 - thin-provisioned volume 56
- repairsevdiskcopy command 755
- repairvdiskcopy command 756
- rescuenode command 618
- resetleds command 411
- resetpassword command 618
- restartservice command 619
- restore command 133, 135
 - help 129
- restore commands
 - clear 131
 - restore 133, 135
- restorevolume
 - volume commands 757
- rmarray command 124
- rmcloudaccount
 - clustered system commands 150
- rmddnsrserver
 - Email and event notification
 - commands 327
- rmemailserver command 371
- rmemailuser command 371
- rmfconsistgrp command 447

- rmfcmap command 448
- rmhost command 476
- rmhostcluster command
 - host commands
 - rmhostcluster 477
- rmhostclustermember command
 - host commands
 - rmhostclustermember 478
- rmhostiogrp command 480
- rmhostport command 481
- rmiscsistorageport command 285
- rmkeyserver
 - clustered system commands 491
- rmdapserver command 648
- rmmdisk command 531
- rmmdiskgrp command 532
- rmmetadatavdisk command 761
- rmnode / rmnodecanister command 285
- rmpartnership command 566
- rmportip command 288
- rmrcconsistgrp command 567
- rmrcrelationship command 567
- rmsnmpserver command 372
- rmsyslogserver command 327
- rmthrottle command 289
- rmuser command 648
- rmusergrp command 649
- rmvdisk command 759
- rmvdiskaccess command 763
- rmvdiskcopy command 762
- rmvdiskhostmap command 763
- rmvolume command 768
- rmvolumebackupgeneration
 - volume commands 765
- rmvolumeecopy command 770
- rmvolumegroup
 - volume commands 764
- rmvolumehostclustermap command
 - host commands
 - rmvolumehostclustermap 479
- running
 - PuTTY plink utility 4

S

- SAN Volume Controller
 - front panel password 14
 - properties 20
- SAN Volume Controller library
 - related publications xi
- scanning
 - Fibre Channel network 20
 - rebalancing MDisk access 20
- secure shell
 - PuTTY 3
- secure shell (SSH)
 - authenticating logins 1
 - client
 - AIX or Linux 6
 - Windows 2
 - creating keys 2
 - overview 1
- secure shell client
 - preparing for CLI on AIX 6
 - preparing for CLI on Linux 7
- security 1

- sending
 - comments xiii
- sendinventoryemail command 372
- service commands
 - chnodeled 605
 - chserviceip 607
 - chwwnn 611
 - installsoftware 613
 - leavecluster 614
 - lscmdstatus 589
 - lsfcportcandidate 200
 - lsfiles 590
 - lssasportcandidate 258
 - lsservicenodes 593
 - lsservicerecommendation 594
 - lsservicestatus 315, 595
 - metadata 615
 - rescuenode 618
 - resetpassword 618
 - restartservice 619
 - setpacedccu 621
 - settempsshkey 621
 - snap 622
 - startservice 623
 - stopnode 623
 - stopservice 624
 - t3recovery 625
- service information commands 587
- Service information commands
 - activatefeature 419
 - deactivatefeature 422
 - lsfeature 423
- service mode
 - commands 601
 - information commands 603
- service mode commands
 - applysoftware 601
 - cleardumps 601
 - dumperrlog 601
 - exit 603
 - overview 601
- service mode information commands
 - ls2145dumps 603
 - lscimomdumps 603
 - lsclustervpd 603
 - lserrlogdumps 603
 - lsfeaturedumps 603
 - lsiostatsdumps 603
 - lsiotracedumps 603
 - lsmdiskdumps 603
 - lssoftwaredumps 605
 - overview 603
- service task commands 605
 - cpfiles 612
- Service task commands
 - help 194
- setlocale command 328, 620
- setpacedccu command 621
- setpwdreset command 290
- setquorum command 516
- setsystemtime command 289
- settempsshkey command 621
- settimezone command 291
- setting
 - quorum disks 25
- settings
 - email server 82

- settings (*continued*)
 - error notification 80
 - event notification 79
- shortcut keys
 - keyboard 945
- showtimezone command 291
- shrinkvdiskspace command 62, 766
- snap command 622
- SNMP traps 79
- software
 - updating using the command-line interface (CLI) 85
- software packages
 - listing 605
 - viewing 321
- splitvdiskcopy command 772
- SSH (secure shell)
 - client system
 - preparing to issue CLI commands 6
- SSH *See* secure shell 1
- SSH keys
 - creating 2
- startemail command 373
- startfcconsistgrp command 449
- startfcmap command 450
- startrcconsistgrp command 568
- startrcrelationship command 571
- startservice command 614, 623
- startstats command 292
- statistics 275, 630, 633, 637, 638, 643, 648, 649
- stopcluster command 293
- stopemail command 374
- stopfcconsistgrp command 452
- stopfcmap command 453
- stopnode command 623
- stopping
 - FlashCopy mapping 35
- stoprcconsistgrp command 574
- stoprcrelationship command 576
- stopservice command 624
- stopstats command 293
- stopsystem command 294
- storage pool commands
 - addmdisk 497
 - chmdiskgrp 517
 - mkmdiskgrp 527
 - overview 517
 - rmmdisk 531
 - rmmdiskgrp 532
- storage pools
 - creating using the CLI 22
- subnet mask
 - changing 72
- svcconfig command 129
- svqueryclock command 329
- switchrcconsistgrp command 578
- switchrcrelationship command 579
- SYSLOG 80
- system
 - recovering nodes 57
- system log
 - information 80
- systems
 - adding nodes 51
 - deleting nodes 68

- systems (*continued*)
 - gateway address
 - changing 72
 - removing nodes 68

T

- t3recovery command 625
- testcloudaccount
 - clustered system commands 150
- testemail command 374
- testkeyserver
 - clustered system commands 492
- testldapserver command 649
- time
 - setting clustered system
 - using the CLI (command-line interface) 11
- time zones 257
- traceroute
 - clustered system commands 600
- trademarks 949
- triggerenclosedump command 412
- triggerlivedump command 495
- triggermddiskdump command 517

U

- updating
 - software using the command-line interface (CLI) 85
- Updating
 - license
 - using the CLI (command-line interface) 12
- user groups
 - changing 77
 - modifying 77
- user management commands 627
 - chauthservice 627
 - chcurrentuser 630
 - chuser 635
 - chusergrp 636
 - mkuser 644
 - mkusergrp 645
 - rmuser 648
 - rmusergrp 649
- users
 - creating 8
 - creating using CLI 78
 - modifying using CLI 78

V

- validating
 - volume copies 54
- viewing
 - clustered systems 263
 - Global Mirror
 - consistency groups 549
 - relationships 552
 - I/O groups 206
 - Metro Mirror
 - consistency groups 549
 - relationships 552

- Viewing
 - license
 - using the CLI (command-line interface) 12
- vital product data (VPD)
 - listing 603
 - viewing 231
- volume
 - copying 655
 - creating 28
 - deleting a copy 31
 - determining mappings 50
 - expanding 60, 61
 - managed disks (MDisks)
 - relationships 50
 - MDisks (managed disks)
 - relationships 50
 - migrating 66
 - shrinkvdiskspace command 62
 - viewing FlashCopy mappings 716
- volume commands
 - addvdiskcopy 655
 - addvolume copy 653
 - analyzevdisk 663
 - analyzevdiskbysystem 663
 - backupvolume 664
 - backupvolume group 665
 - chvdisk 666
 - chvolume group 670
 - expandvdiskspace 671
 - lscontrollerdependentvdisks 336
 - lsmetadataavdisk 677
 - lsrepairsevdiskcopyprogress 678
 - lsrepairvdiskcopyprogress 680
 - lssevdiskcopy 681
 - lsvdiskanalysis 703
 - lsvdiskanalysisprogress 706
 - lsvdiskcopy 707
 - lsvdisksyncprogress 722
 - lsvolumebackup 723
 - lsvolumebackupgeneration 725
 - lsvolumebackupprogress 727
 - lsvolume group 729
 - lsvolumerestoreprogress 731
 - mkimagevolume command 749
 - mkmetadataavdisk 733
 - mkvdisk 734
 - mkvdiskhostmap 743
 - mkvolume command 745
 - mkvolume group 748
 - overview 653
 - recovervdisk 753
 - recovervdiskbyiogrp 754
 - recovervdiskbysystem 754
 - repairsevdiskcopy 755
 - repairvdiskcopy 756
 - restorevolume 757
 - rmmetadataavdisk 761
 - rmvdisk 759
 - rmvdiskcopy 762
 - rmvdiskhostmap 763
 - rmvolume command 768
 - rmvolumebackupgeneration 765
 - rmvolume copy command 770
 - rmvolume group 764
 - shrinkvdiskspace 766
 - splitvdiskcopy 772

- volume copies
 - validating 54
- volume disks
 - removing 762
- volume extent
 - viewing 714
- Volume Mirroring
 - memory 26
- volume) 61
- volumes
 - adding a copy 31
 - converting
 - from image mode to managed mode 67
 - creating 734
 - determining name of 49
 - listing node dependent 48
 - recovering 58
 - recovering from offline
 - using CLI 56
 - using the CLI 58
 - viewing 687
 - viewing disks 718

W

- writesernum command 329



Printed in USA

GC27-2287-14

